MIT Technology Review

VOL. 120 NO. 3 MAY/JUNE 2017 US $6.99/CAN $7.99

Feature p. 42

A 3-D Printer That Really Matters

Feature p. 78

Cancer Cures For a Lucky Few

Feature p. 28

Time to Consider Geoengineering?

Mysterious Machines

Artificial intelligence is a black box that thinks in ways we don't understand. That's thrilling and scary. p. 54

From the Editor

Make America Great Again

The president's proposed cuts to research funding would cripple American innovation. We should be spending more on R&D, not less.

"Most things may never happen," Philip Larkin wrote, and presidential budget requests never pass unaltered, even when one party controls the executive and legislative branches of government. Congress is jealous of its power to appropriate, and legislators fight to protect programs important to their districts. "America First: A Budget Blueprint to Make America Great Again," released in March, isn't even a presidential budget request. It's a sketch, lacking details, of the priorities of the administration of Donald Trump, and an opening bid in the art of the deal. But the blueprint eloquently expresses the things the White House cares about, and federal funding of science and technology isn't among them.

The blueprint proposes that the National Institutes of Health should lose $6 billion, or over 18 percent of its budget, and the Department of Energy's Office of Science $900 million, or nearly 20 percent of its funding. The White House wants a 40 percent reduction in science programs at the Environmental Protection Agency, and a 26 percent reduction in research at the National Oceanic and Atmospheric Administration. Oddly, the National Science Foundation and its $7.5 billion budget aren't mentioned, nor are the Defense Advanced Research Projects Agency (DARPA) and its nearly $3 billion budget. Perhaps the formal presidential budget request, due in May, will be more specific. But the White House is clear enough about the Advanced Research Projects Agency–Energy (ARPA-E), a small agency with $300 million in appropriations that has developed energy technologies such as advanced batteries: it should be eliminated, because "the private sector is better positioned to finance disruptive energy research and development and to commercialize innovative technologies."

Whence this passion for culling science and technology budgets? Some part must be bluntly mathematical: the blueprint vows to increase defense spending by $54 billion while leaving Social Security and Medicare untouched, all without increasing the federal deficit. Something had to give, and the NIH was the loser. The punishment delivered to the EPA and NOAA doubtless derives from the administration's derisive attitude toward climate science (which Mick Mulvaney, the director of the Office of Management and Budget, has called a "waste of your money" and the president has called a "hoax"). But why eliminate ARPA-E and reduce DOE's spending by so much?

Trump's best policy would be to return funding of R&D to the scale it enjoyed during the 1960s, when the U.S. spent more than the rest of the world combined.

To a remarkable degree, the White House's blueprint cribs from a "Blueprint for Balance" published by the conservative Heritage Foundation last year. Heritage's plan also calls for the elimination of ARPA-E, and in strikingly similar language. It castigates the DOE because its

GUIDO VITTI

programs on basic energy science "stray from fundamental research into commercialization," and it demands that the government eliminate such programs because "private companies are capable of fulfilling these roles, whether through their own laboratories or by funding university research." Throughout, the text breathes fundamental misunderstandings about how scientific breakthroughs become commercial technologies.

Every year since 2002, *MIT Technology Review* has selected the 10 breakthrough technologies we think most likely to have a significant impact. Gene therapy (2017), the cell atlas (2017), electronic interfaces to reverse paralysis (2017), immune engineering (2016), liquid biopsies (2015), and brain organoids (2015) were all the products of years of expensive, risky research that relied on the U.S. government for funding because there was no clear economic incentive for industry to undertake the work. That research helped push the technologies into the marketplace. But pulling breakthrough technologies into commercial use often requires still more research: very efficient solar thermophotovoltaic cells (2017), precise gene editing in plants (2016), and supercharged photosynthesis (2015) will demand more federally funded development to prove they can scale to industrial use. Companies won't subsidize all of that cost.

Even if we stipulate that the blueprint's demands are the pitch in a complex negotiation with Congress, cuts half so severe would cripple the American method of innovation inaugurated by Vannevar Bush, FDR's science advisor: funding agencies distribute competitive grants to outside researchers at universities, and their research, if promising, is turned into intellectual property by technology transfer offices, invested in by venture capitalists, and commercialized by entrepreneurs. For 70 years, Bush's methodology was the primary driver of American productivity and competitiveness. It created Silicon Valley and the regions that imitate it, and it ushered in the American Century.

For all its payoffs in domestic jobs and prosperity, the U.S. government's historical investment in R&D was undertaken in an optimistic impulse of generosity.

If President Trump truly wishes to make America great again, his best policy would be to return funding of research and development to the scale it enjoyed during the 1960s, when the U.S. spent more on R&D than the rest of the world combined. During the mid-1960s, the federal government's spending on R&D reached a peak of 1.9 percent of the gross domestic product, roughly 10 percent of total outlays, according to the U.S. Budget's Historical Tables; in 2016, the Obama administration estimated that the government would spend only 0.7 percent of GDP on R&D, or 3.4 percent of total outlays. The nation should be spending more, not less, on science and technology.

"The drastic reductions in funding for science agencies in the president's budget would put a tourniquet on the flow of ideas that produce new companies and new jobs," says Susan Hockfield, the president of the American Association for the Advancement of Science (and the previous president of MIT). "Federally sponsored research spawns new businesses and new industries, returning $40 to our economy on every research dollar. Federal investment in R&D has driven roughly half of America's economic growth. All of the technologies we enjoy today—cell phones, GPS, vaccines, new cancer cures—use discoveries from federally funded research."

There is a final, broader point, which will be unpopular in the era of America First: for all its payoffs in domestic jobs and prosperity, the U.S. government's historical investment in R&D was undertaken in an optimistic impulse of generosity. It was understood to be part of our larger duties and responsibilities to the world. Ronald Reagan, in a weekly radio address in 1988, said, "The remarkable thing is that although basic research does not begin with a particular practical goal, when you look at the results over the years, it ends up being one of the most practical things government does." Science and technology can be neutral, evidence-based pursuits that take the heat out of partisan battles, and they can solve big problems, grow wealth, and expand human possibilities.

But write to me at jason.pontin@technologyreview.com and tell me what you think.

Contents

May/June 2017

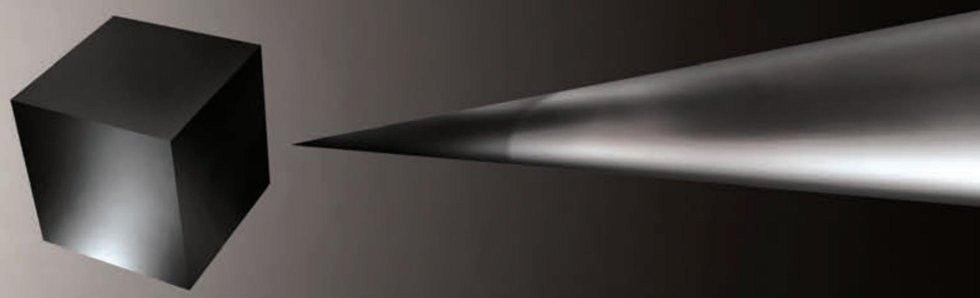

ON THE COVER

Illustration by Keith Rankin

KEITH RANKIN

Be a Part of the **Solution**

Solve is a community of leaders and change makers committed to solving the world's most pressing challenges through open innovation and partnership.

Starting May 10, 2017, we're accepting solutions to challenges in four categories:

- **Sustainability**
- **Health**
- **Economic Prosperity**
- **Education**

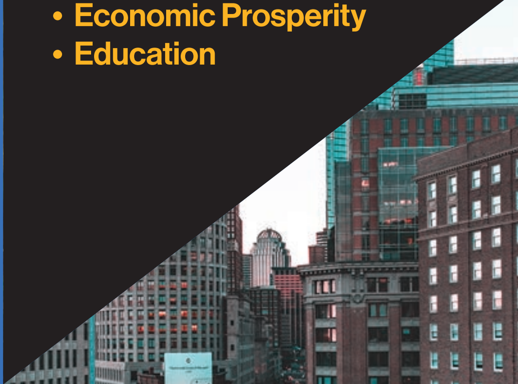

Submit a solution:
solve.mit.edu/challenges

For membership inquiries:
membership@solve.mit.edu

SOLVE

CUSTOMER SERVICE AND SUBSCRIPTION INQUIRIES

National: 800-877-5230 / International: 903-636-1115 / E-mail: technologyreview@pubservice.com /Web: www.technologyreview.com/customerservice / MIT Records: 617-253-8270 (alums only) Reprints: techreview@wrightsmedia.com / 877-652-5295, Licensing and permissions: licensing@technologyreview.com

Technology Review
One Main Street, 13th Floor
Cambridge, MA 02142
Tel: 617-475-8000

The mission of *MIT Technology Review* is to equip its audiences with the intelligence to understand and contribute to a world shaped by technology.

Technology Review, Inc., is an independent nonprofit 501(c)(3) corporation wholly owned by MIT; the views expressed in our publications and at our events are not always shared by the Institute.

De Technologia non multum scimus. Scimus autem, quid nobis placeat.


THINK BIG
TAKE RISKS
INNOVATE!
R&D Funding Program
The National Reconnaissance Office Director's Innovation Initiative (DII) Program funds cutting-edge scientific research in a high-risk, high-payoff environment to discover innovative concepts and creative ideas that transform overhead intelligence capabilities and systems for future national security intelligence needs. The program seeks out the brightest minds and breakthrough technologies from industry, academia, national laboratories, and U.S. government agencies.
Visit the website for program history, frequently asked questions, proposal guidance, and Broad Agency Announcement and Government Sources Sought Announcement requirements.
703.808.2769
NATIONAL RECONNAISSANCE OFFICE
UNITED STATES OF AMERICA
https://acq.westfields.net

Feedback

E-mail letters@technologyreview.com

Write *MIT Technology Review*
One Main Street, 13th Floor
Cambridge, MA 02142

Please include your address, telephone number, and e-mail address.

Letters and comments may be edited for both clarity and length.

Letters and Comments

MIT Technology Review
Volume 120, Number 2

Let's Not Get Hysterical

There is no doubt a lot of nonsense and hysteria in the air about the topic of automation. With no consensus in place yet on what is even happening, a serious need remains for technologically alert, humanitarian reflection on the relevant trends and what should be done about them. In that spirit, David Rotman's sober essay—"The Relentless Pace of Automation" (March/April 2017)—performed a real service. Without succumbing to doomsaying, Rotman's cogent review of new developments and emerging issues raised important questions in a spirit of prudence. Most notably, Rotman makes clear what is glaringly obvious: that policymakers and political leaders are unprepared to deal with the ways technology may be reshaping employment, including by displacing workers in a variety of occupations. Ultimately, Rotman's point seems spot on: without getting hysterical, we should not be complacent about the arrival of truly brilliant technologies.

Mark Muro is a senior fellow and the director of policy at the Metropolitan Policy Program at the Brookings Institution.

Policymakers and political leaders are unprepared to deal with the ways technology may be reshaping employment.

Get Ready for the Robots

Regarding David Rotman's piece on "The Relentless Pace of Automation": A little while back I was at a small gathering that included many of the rock stars of artificial intelligence and machine learning. I gave a talk to the group, relaying how mainstream economists have thought about the effects that tech progress has on jobs and wages. The reason we have what's labeled a "Luddite fallacy" is that for most of the 200-plus years of the industrial era, tech progress has created many more jobs than it's destroyed, and has steadily lifted wages for people at all levels of skill and education. I also showed evidence, though, that job and especially wage growth have slowed down recently in the rich world, and that the middle class in industrial countries has been getting hollowed out.

So the question is—is this time finally different? Will the digital technologies of today and tomorrow have a different effect on jobs and wages? I concluded by stating what I believe, which is that it's a distinct possibility that we should take seriously, but that no one knows for sure.

The reaction to my talk was fascinating to me. Several times during the rest of the conference someone would come up to me during a break and essentially say, "I think you're underestimating what's going on. This time *is* different, and we'd better get ready."

Andrew McAfee is co-director of the MIT Initiative on the Digital Economy and the associate director of the Center for Digital Business at the MIT Sloan School of Management.

David Rotman responds:

Both Mark Muro and Andrew McAfee have done important work in helping us better understand the complex effects of automation and artificial intelligence on jobs. There are still a lot of unknowns, including how fast and how radical the technological changes will be and how adeptly the labor force will adapt. But it is clear our politicians need to be far more aware of the ongoing distress that automation is causing many workers. As Muro has recently pointed out, President Trump's budget outline (see "Make America Great Again," page 2), issued in March, proposed eliminating just the kinds of programs, such as the Manufacturing Extension Partnership, that companies and workers will need to better prepare for the changes. Helping those already hurt by the forces of automation shouldn't be a partisan issue.

Follow Us Twitter @techreview
Facebook @technologyreview
Instagram @technologyreview

Twitter Votes

Reactions to our stories on paying with your face, botnets of things, and virtual-reality filmmaking.

If you could use your face to pay for things, would you?

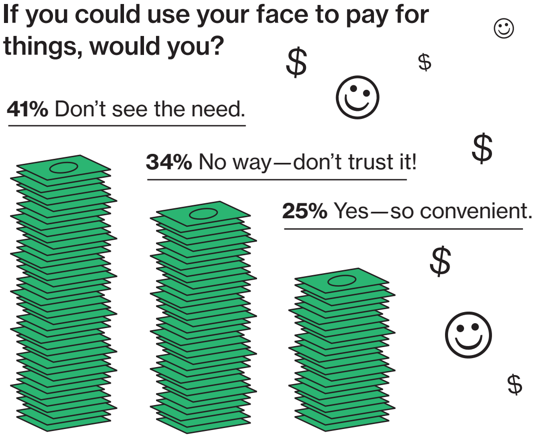

How does the increasing threat of botnets make you feel about devices connected to the Internet of things?

14%	I'll stop buying them.
36%	Too convenient not to use.
50%	Never wanted them anyway.

Do you think people in Hollywood will get the hang of making compelling VR movies?

36%	No, it's not their medium.
42%	Yes, they're creative.
22%	With help from Google.

Views

Janos Pasztor

Elaine Lissner

Jerry Kaplan

ANDY FRIEDMAN

CLIMATE CHANGE

Cooling-Off Period

We have to at least consider geoengineering. And that's where the problems start.

The 2015 Paris climate agreement was a major milestone, but the truth is, achieving its ambitious goal of keeping temperatures to within 1.5 °C to 2 °C of preindustrial levels would require rates of mitigation far in excess of what's been achieved—or even what's been planned.

Because of this, more people are contemplating geoengineering—notably solar radiation management, which involves reflecting a portion of the sun's radiation back into space. The idea raises many questions. We don't know how effective it would be, and we don't fully understand its potential impacts. There are also ethical issues about its use and its governance.

We need to acknowledge that the aggregate environmental and socioeconomic risks of solar radiation management would probably be small in comparison with the benefits of reducing global temperatures. But those benefits and harms would be unequally spread among regions of the world, and between current and future generations.

In the absence of multilateral agreements, there's no way of controlling who might execute such a geoengineering plan. It's possible that a small group of countries, or a single country, or a large company, or even a wealthy individual might take unilateral action on geoengineering. Others might subsequently engage in their own climate engineering strategies to counter such action.

To avoid such a future, we should establish global governance frameworks. Currently there's really only one forum that could give legitimacy to any such framework for geoengineering: the United Nations General Assembly.

Here are the kinds of questions such a framework would have to address: Who controls the "global thermostat"? How would decisions be made to balance the need to reduce the global temperature with the inequality of regional and local impacts across the globe? How would trans-border and trans-generational ethical issues be addressed? How would the required governance frameworks withstand potentially substantial geopolitical changes over the decades, and possibly centuries, over which they would need to be deployed? How might such techniques be deployed without undermining the will to cut emissions (which will continue to be necessary no matter what)? How would decisions relating to the rate of starting, continuing, and stopping those techniques be governed?

This last issue is of particular concern, as suddenly stopping a geoengineering scheme would result in a rapid and probably catastrophic rise in temperatures. Many of these governance issues might turn out to be unresolvable and thus might keep us from attempting this kind of fix in the first place.

The research community has been addressing many of these issues, but the global policy community and the public have not. It's time to begin doing so.

Janos Pasztor is executive director of the Carnegie Climate Geoengineering Governance Initiative.

BIOTECHNOLOGY

The Elusive Male Pill

Many forces have slowed the development of better contraceptives for men.

It seems as if every few months the press heralds a scientific breakthrough that could lead to a new male contraceptive—in five to 10 years. But then the years go

by and the promised contraceptive never appears. Why? What's the holdup?

Maybe one problem is that people have been using the wrong tool for the job. Chemists in the 1950s changed society forever when they figured out how to shut off female ovulation with synthetic hormones. So you'd think we'd just have to repeat the trick with men, right?

Not quite. Many a researcher has lamented that it's harder to stop millions of sperm than it is to stop one egg—a complaint that sounds intuitively logical. But it goes beyond that. Men are not women, and men's anatomy is not women's anatomy. So it doesn't make sense to use the same tools—hormones—that we used for women.

In men, all the sperm swim through one tiny tube, the vas deferens. So rather than shut them off at the source, we should disrupt their transit.

Some of the most intriguing methods in this direction make use of muscular action to clamp down on the vas deferens. My own group is working on something called Vasalgel, a polymer gel that blocks or filters out sperm in the vas deferens tube. A plant-derived compound in advanced trials in Indonesia interferes with sperm's ability to penetrate the egg. These are the elegant approaches we need to be advancing, not hammering the entire body with hormones.

So why haven't we yet finished the job?

For starters, clinical trials are a multimillion-dollar process. If you're a researcher developing one of these methods, where do you turn? Big pharma is not interested—the liability in treating healthy young people for years is high, and if you're already the company selling contraceptives to those men's partners, the payoff is low. Nonprofits try to do the development work but are hampered by shoestring budgets and a lack of business savvy. Major foundations such as the Bill & Melinda Gates Foundation are stepping up to fill the gap for female contraceptive development but have determined that men in their target areas—ultra-poor regions such as Bihar, India, and sub-Saharan Africa—are not ready for male contraception.

Maybe the best option is something called social venture enterprise—turning to social investors who would like to get their money back if it goes well but who are also committed to seeing an affordable product reach the market.

Big pharma is not interested. If you're already selling contraceptives to women, the payoff for a male pill is low.

We have a waiting list of more than 38,000 men and women hoping for news of clinical trials for Vasalgel. There's clearly a market for a male contraceptive. Now we just have to get them one.

Elaine Lissner is the founder of Parsemus Foundation, which aims to advance low-cost evidence-based medicines.

ARTIFICIAL INTELLIGENCE

AI's PR Problem

The spooky-sounding name was probably never a good idea.

Smart people like Bill Gates and Steven Hawking have warned that artificial intelligence could threaten the human race.

And they're not the only ones worried. The Committee on Legal Affairs of the European Parliament recently issued a report calling on the EU to require intelligent robots to be registered, in part so their ethical character can be assessed. The "Stop Killer Robots" movement, opposed to the use of so-called autonomous weapons in war, is influencing both United Nations and U.S. Defense Department policy.

Artificial intelligence, it seems, has a PR problem. While it's true that today's machines can credibly perform many tasks (playing chess, driving cars) that were once reserved for humans, that doesn't mean the machines are growing more intelligent and ambitious. It just means they're doing what we built them to do.

Machines have been taking over skilled work for centuries, but the machines don't aspire to better jobs and higher employment. Jacquard looms replaced expert needleworkers in the 19th century, but they didn't spell doom for tailors. Until the mid-20th century we relied on our best and brightest to do arithmetic, but now that comparably capable devices are given away as promotional trinkets at trade shows, the mathematically minded among us can focus on tasks that require broader skills, like statistical analysis.

I'd suggest that one problem is the name itself. Had artificial intelligence been named something less spooky, it might seem as prosaic as operations research or predictive analytics. Perhaps a less provocative description would be something like "anthropic computing," a broad moniker that could encompass efforts to design biologically inspired computer systems, machines that mimic the human form or abilities, and programs that interact with people in natural, familiar ways.

Yes, we should be careful about how we deploy AI, but not because we are summoning some mythical demon. Instead, we should accept these remarkable inventions for what they really are—potent tools that promise a more prosperous and comfortable future.

Jerry Kaplan teaches at Stanford University. His latest book is Artificial Intelligence: What Everyone Needs to Know, *from Oxford University Press.*

Upfront

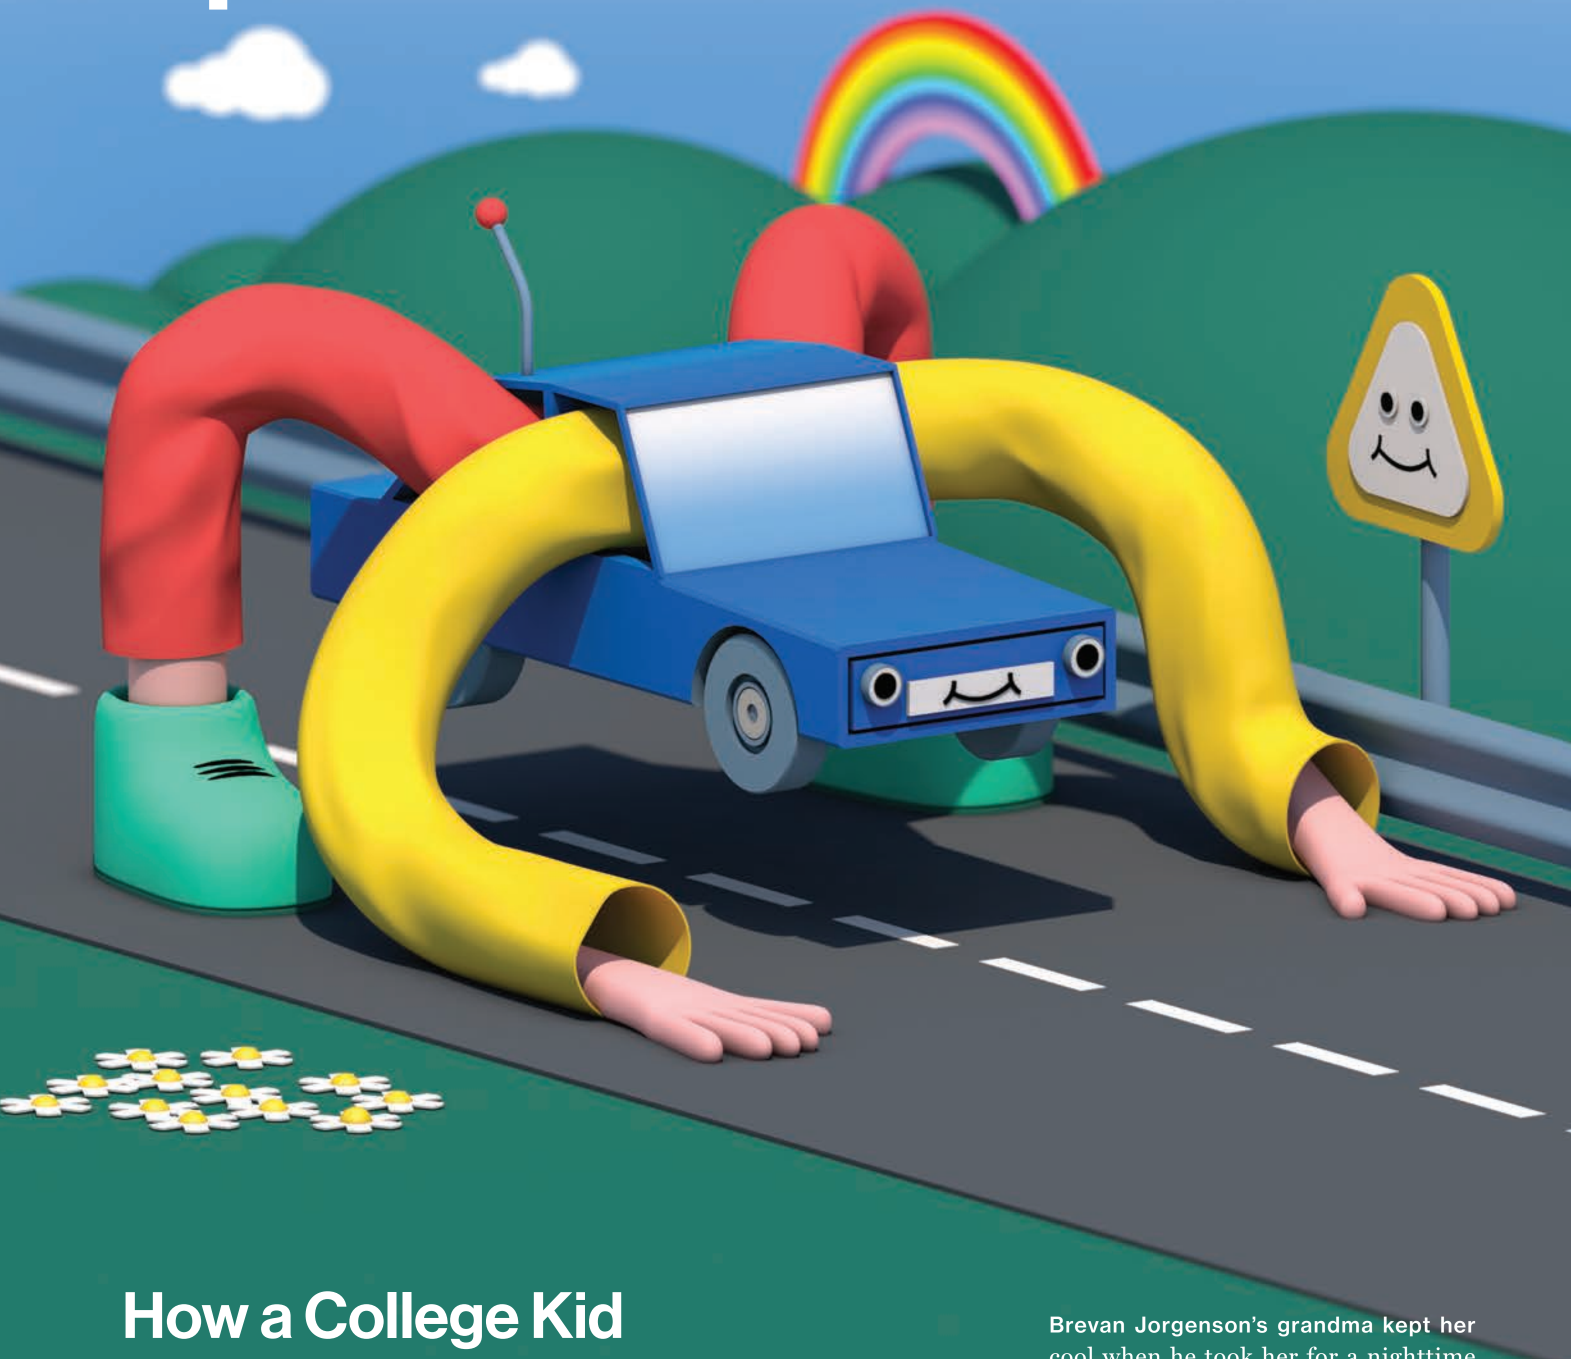

JACK SACHS

How a College Kid Made His Honda Civic Self-Driving for $700

Who needs a Tesla when you can build your own automated copilot using free hardware designs and software available online?

Brevan Jorgenson's grandma kept her cool when he took her for a nighttime spin in the Honda Civic he's modified to drive itself on the highway. A homemade device in place of the rearview mirror can control the brakes, accelerator, and steering, and it uses a camera to identify road markings and other cars.

Upfront

"She wasn't really flabbergasted—I think because she's seen so much from technology by now," says Jorgenson, a senior at the University of Nebraska, Omaha. Others are warier of the system, which he built using plans and software downloaded from the Internet, plus about $700 in parts. Jorgenson says the fact that he closely supervises his home-brew autopilot hasn't convinced his girlfriend to trust the gadget's driving. "She's worried it's going to crash the car," he says.

Many tech and auto companies have begun testing modified cars on the road in recent years. Jorgenson's vehicle is in the vanguard of a more ragged grassroots test fleet taking shape as tinkerers around the world strive to upgrade their own vehicles with computing gear that can share driving duties.

Motivation comes from the fun and challenge of getting the technology working—and the prospect of making driving easier. Kiki Jewell, who set out to make her Chevy Bolt self-driving as a learning exercise, says her spouse has been strongly supportive, partly out of self-interest. "My husband's happy I'm interested to ease his commute," says Jewell, who lives in the Bay Area.

Jewell and Jorgenson's projects were enabled by a fit of pique last October by the founder of Comma.ai, a San Francisco startup, which was developing a $999 device that could upgrade certain vehicles to steer themselves on the highway and follow stop-and-go traffic. Founder George Hotz abruptly cancelled plans to launch the product after receiving a letter from the National Highway Traffic Safety Administration asking questions about its functionality.

Last November, he released the company's hardware designs and software for free, saying he wanted to empower researchers and hobbyists. (Hotz didn't respond to requests to talk about his strategy.)

Jorgenson set about ordering the parts needed to build Comma's device, the Neo, the same day Hotz dumped the plans online. He had been following Comma's fortunes, and he happened to own a 2016 Honda Civic, one of the two models supported by the company's software (the other is the 2016 Acura ILX).

A Neo is built from a OnePlus 3 smartphone equipped with Comma's now-free Openpilot software, a circuit board that connects the device to the car's electronics, and a 3-D-printed case. Jorgenson got the case printed by an online service and soldered the board together himself.

He first put his life in the device's hands in late January after an evening college class. "It was dark on the interstate, and I tested it by myself because I figured if anything went wrong I didn't want anybody else in the car," says Jorgenson. "It worked phenomenally."

Consumers have significant flexibility in legally changing their own cars.

Subsequent tests revealed that the Neo would inexplicably pull to the right sometimes, but a software update released by Comma quickly fixed that. Now fully working, the system is similar in capabilities to the initial version of Tesla's AutoPilot.

Comma's plans and software aren't the only resources out there for wanna-be self-driving-car builders. Neodriven, a startup based in Los Angeles, recently started selling a pre-built Neo device that works with Comma's Openpilot; it costs $1,495. Online-education platform Udacity has released code used in its autonomous-car research program, and students in one of its courses are actively improving and expanding it.

Bryant Walker Smith, a law professor at the University of South Carolina, says that federal and state laws probably don't pose much of a barrier to those with a desire to upgrade their vehicle to share driving duties. NHTSA has authority over companies selling vehicles and systems used to modify them, but consumers have significant flexibility in making changes to their own vehicles, says Smith, who advises the U.S. Department of Transportation on legal issues and automation.

Anyone using a home-built Neo will still have to comply with state rules requiring responsible driving, though. (Comma's Openpilot software tries to help with that: it complains if the driver doesn't touch the wheel every five minutes, and it asks for human intervention if it's having trouble interpreting the road ahead.) And in the event of a crash, using a home-built driving aid might raise eyebrows. "Just because you can legally operate it doesn't mean you are not civilly liable," says Smith.

Ariel Núñez, a software developer in Barranquilla, Colombia, hopes that the work of hobbyists like himself will show how existing cars could be made significantly safer—an alternative vision to the one chased by giant companies that are focused on ending the need for human drivers. Núñez is using code from Comma and Udacity to try to get his Ford Fusion to automatically slow down when it sees traffic signs, speed bumps, or potholes. He hasn't tested it on the road but has got the accelerator and steering control working, and he had a near miss with a tree.

"I am less interested in full autonomy and more in preventing rear-endings," he says. "A lot of existing cars can be retrofitted." *—Tom Simonite*

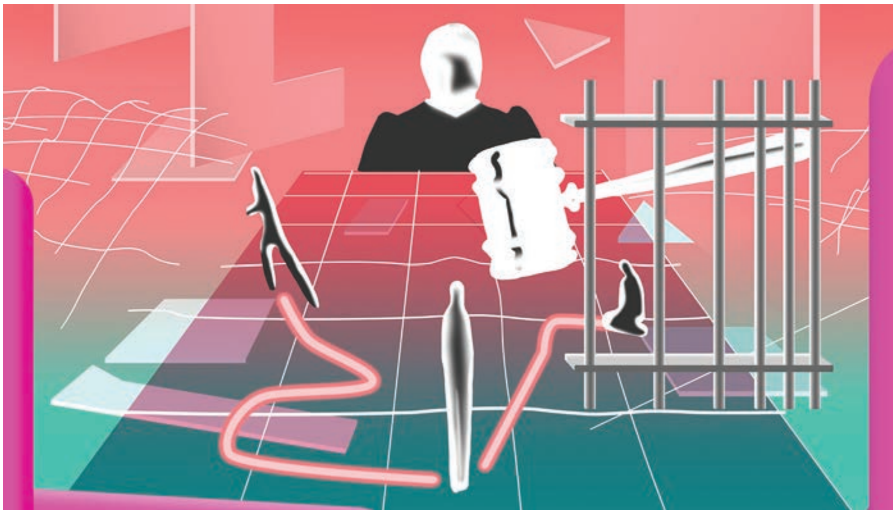

How Software Could Help Judges Reduce Crime

An algorithm can predict what defendants might do if released while awaiting trial.

When should a criminal defendant be required to await trial in jail rather than at home? Software could significantly improve judges' ability to make that call. In a study from the National Bureau of Economic Research, economists and computer scientists using data from hundreds of thousands of cases in New York City trained an algorithm to predict, from rap sheets and court records, whether defendants were flight risks. When tested on over 100,000 cases that it hadn't seen before, the algorithm proved better than judges at predicting what defendants would do after release. Jon Kleinberg, a computer science professor at Cornell who was involved in the research, says one goal of the project was to show policymakers the potential benefits to society of using machine learning in the criminal justice system. "This shows how machine learning can help even in contexts where there's considerable human expertise being brought to bear," says Kleinberg.

The researchers estimate that for New York City, their algorithm's advice could cut crime by defendants awaiting trial by as much as 25 percent without changing the number of people waiting in jail. Alternatively, it could be used to reduce the jail population awaiting trial by more than 40 percent, while leaving the crime rate by defendants unchanged. As a bonus, gains like those were possible while reducing the proportion of African-Americans and Hispanics in the jail population.

The algorithm assigns defendants a risk score based on data pulled from records for their current case and their rap sheet. The only demographic data it uses is age—not race or anything else. Kleinberg says the algorithm could work as an assistant to judges that suggests they reconsider when they are about to approve release of a suspect very likely to fail to show up in court, or to commit a crime while awaiting trial.

—Tom Simonite

TO MARKET

Here One

Wireless Smart Earbuds

COMPANY:
Doppler Labs

PRICE:
$299.99

AVAILABILITY:
Now

Your earbuds can probably play music, take phone calls, and summon a personal assistant like Siri or Google Now, but they're not very smart. Startup Doppler Labs is one of several companies that have figured this out. Its wireless earbuds let you manipulate the world around you by filtering all kinds of sounds (turn down the din in a restaurant, for instance) and helping you home in on the ones you want to enhance (turn up the person speaking to you). The ear gear, called Here One, is too pricey for the average consumer and still has plenty of kinks to work out. Still, its capabilities offer a fascinating look at where wearable technology is heading. —*Rachel Metz*

JON HAN; COURTESY OF DOPPLER LABS INC.

Upfront

AI Software Learns to Make AI Software

Machine-learning experts are in short supply. Could software do their jobs instead?

Progress in artificial intelligence causes some people to worry that software will take jobs such as driving trucks away from humans. Now leading researchers are finding they can make software that learns to do one of the trickiest parts of their own jobs—designing machine-learning software.

In one experiment, researchers at the Google Brain artificial-intelligence research group had software design a machine-learning system to take a test used to benchmark software that processes language. What it came up with surpassed previously published results from software designed by humans. Several other groups have also reported progress on getting learning software to make learning software. They include researchers at the nonprofit research institute OpenAI (which was cofounded by Elon Musk); MIT; the University of California, Berkeley; and Google's other artificial-intelligence research group, DeepMind. If self-starting AI techniques become practical, they could increase the pace at which machine-learning software is implemented across the economy. Companies must currently pay a premium for machine-learning experts, who are in short supply.

One set of experiments from Google's DeepMind group suggests that what researchers are terming "learning to learn" could also help reduce the vast amount of data on a specific task that machine-learning software needs to consume in order to perform it well. The researchers challenged their software to create learning systems for collections of different but related problems, such as navigating mazes. It came up with designs that showed an ability to generalize and pick up new tasks with less additional training than would be usual. The idea of creating software that learns to learn has been around for a while, but previous experiments didn't produce results that rivaled human achievements.

Easing the burden on the data scientist could be a big payoff for companies.

Yoshua Bengio, a professor at the University of Montreal who explored the idea in the 1990s, says that greater computing power and the advent of a technique called deep learning, which has sparked recent excitement about AI, are what's making the approach work. But he notes that so far it requires such extreme computing power that it's not yet practical to think about partially replacing machine-learning experts. Google Brain's researchers describe using 800 high-powered graphics processors to power software that came up with designs for image recognition systems rivaling the best designed by humans.

Otkrist Gupta, a researcher at the MIT Media Lab, believes that will change. He and his colleagues plan to open-source the software behind their own experiments, in which learning software designed deep-learning systems that matched human-crafted ones in object recognition. He thinks companies are well motivated to find ways to make automated machine learning practical. He says, "Easing the burden on the data scientist is a big payoff."

—Tom Simonite

TO MARKET

Cobalt Robot

Security Robot

COMPANY:
Cobalt

PRICE:
N/A

AVAILABILITY:
Now

COURTESY OF COBALT

Would you feel more comfortable interacting with a robot security guard if it looked more like a fancy fabric-covered sculpture than a sterile droid? That's the thinking behind a startup called Cobalt, whose new robot security guards were designed by well-known industrial designer Yves Behar and his company, Fuseproject. Meant for patrolling inside high-end offices and interacting with people, Cobalt's bots have a touch screen on one side that lets office workers communicate with a remote operator when needed. "For the most part, it's providing peace of mind to employees," says Travis Deyle, Cobalt's cofounder. *—Rachel Metz*

Upfront

A Mind-Reading Device Gives Words to "Locked In" Patients

A brain-computer interface records yes and no answers in patients who lack any voluntary muscle movement.

DADU SHIN

In 1995, Jean-Dominique Bauby suffered a massive stroke that left him paralyzed and speechless, with only the ability to blink his left eyelid. Using just that eye, he silently dictated his memoir, *The Diving Bell and the Butterfly*, later adapted into a film. Bauby suffered from "locked-in syndrome," in which patients are completely paralyzed except for some eye movement. Some patients eventually lose even the ability to blink, cutting off all contact with the world and raising questions about whether they are still fully conscious—and, if so, whether they still wish to live.

Now researchers in Europe say they've found out the answer after using a brain-computer interface to communicate with four people who had lost all voluntary movement as a result of Lou Gehrig's disease, or amyotrophic lateral sclerosis.

In response to the statement "I love to live," three of the four replied yes. They also said yes when asked, "Are you happy?" The fourth patient, a 23-year-old woman, wasn't asked these questions because her parents feared she was in a fragile emotional state.

Designed by neuroscientist Niels Birbaumer, now at the Wyss Center for Bio and Neuroengineering in Geneva, the brain-computer interface fits on a person's head like a swimming cap. It measures changes in electrical waves emanating from the brain, as well as blood flow, using a technique known as near-infrared spectroscopy.

To verify that the patients could communicate, Birbaumer's team asked them, over the course of about 10 days of testing, to respond yes or no to statements such as "You were born in Berlin" or "Paris is the capital of Germany" by modulating their thoughts and altering the blood-flow pattern. The answers relayed through the system were consistent about 70 percent of the time, substantially better than chance.

For the patients' family members, Birbaumer says, "the relief was enormous." They were able to communicate with their loved ones after as long as four years of total silence, and to learn they wished to remain alive on ventilators. The team detailed their experiments in a study appearing in the journal *PLOS Biology*.

In 2010, British neuroscientist Adrian Owen first reported that changes in blood flow in certain parts of the brain showed that a patient previously written off as being in a vegetative state was actually conscious.

No one has a clear idea of how many locked-in patients there are, says Jane Huggins, who runs the Direct Brain Interface Laboratory at the University of Michigan, although an estimate by Dutch researchers puts it at less than one in 150,000 people in that country.

Some may be misdiagnosed as comatose because they lack eye movement or it's subtle. Birbaumer and his team say their system could be used as a diagnostic to determine who actually remains conscious and aware, and he hopes to develop a technology to allow people with complete locked-in syndrome to select letters so they can communicate beyond answering yes-or-no questions.

—Emily Mullin

Upfront

Drone Hunters

The French army is training eagles to take down small drones flying in forbidden airspace. Their training begins before birth—the military mounts eggs atop small drones and keeps the young chicks there after hatching and during their early feeding period. The birds are seen as a safer alternative to shooting down the unmanned aircraft in urban areas.

Photograph by Georges Gobet

Upfront

Goldman Sachs Embraces Automation, Leaving Many Behind

Software that works on Wall Street is changing how business is done and who profits from it.

At its height back in 2000, the U.S. cash equities trading desk at Goldman Sachs's New York headquarters employed 600 traders, buying and selling stock on the orders of the investment bank's large clients. Today there are just two equity traders left.

Automated trading programs have taken over the rest of the work, supported by 200 computer engineers. Marty Chavez, the company's chief financial officer, explained all this to attendees at a symposium on computing's impact on economic activity held by Harvard's Institute for Applied Computational Science in January.

The experience of its New York traders is just one early example of a transformation of Goldman Sachs, and increasingly other Wall Street firms, that began with the rise in computerized trading but has accelerated over the past five years, moving into more fields of finance that humans once dominated. Chavez, who was formerly the chief information officer at Goldman, says areas of trading like currencies and even parts of business lines like investment banking are moving in the same automated direction that equities

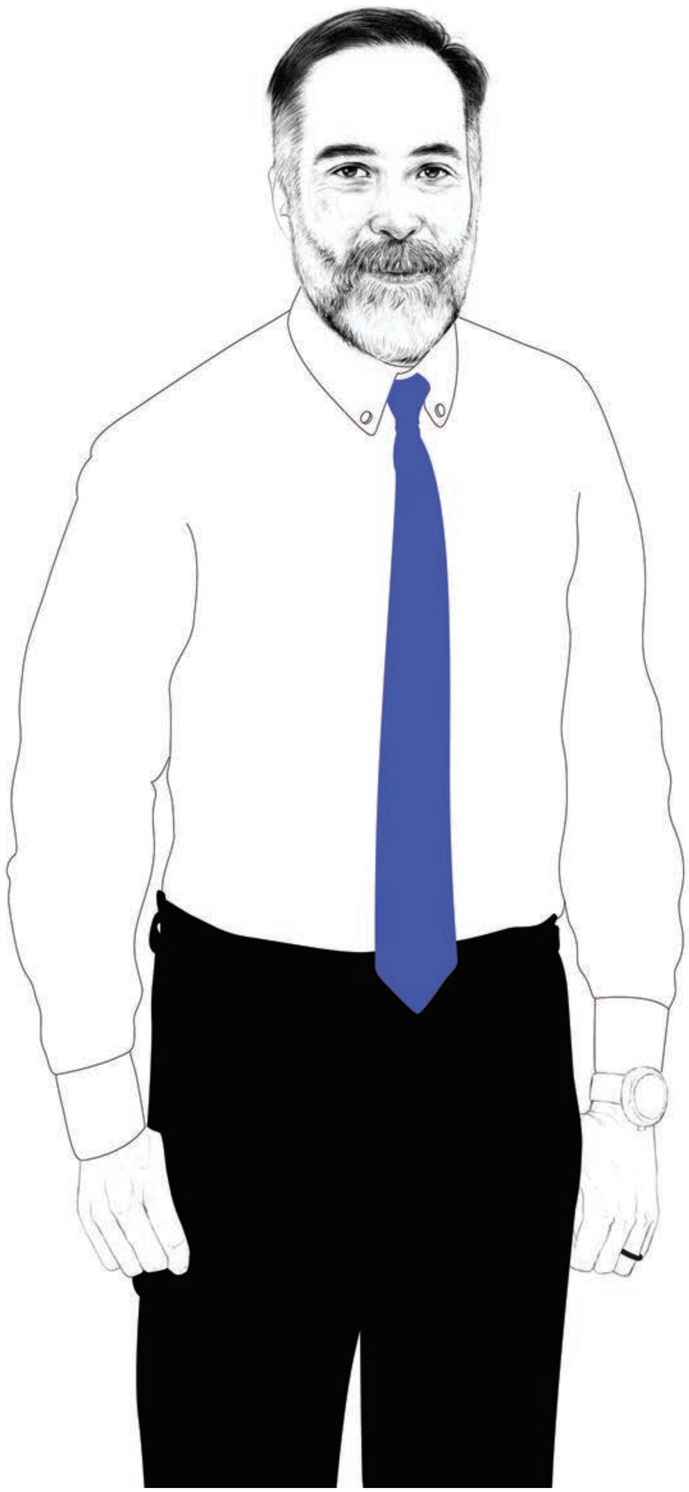

GRAFILU

have already traveled. Today, nearly 45 percent of the revenue from cash equities trading comes from electronic trades, according to Coalition, a U.K. research firm that tracks the industry. In addition to back-office clerical workers, on Wall Street machines are replacing a lot of highly paid people, too.

Average compensation for staff in equities sales, trading, and research at the 12 largest global investment banks, of which Goldman is one, is $500,000 in salary and bonus, according to Coalition. Seventy-five percent of Wall Street compensation goes to these highly paid "front office" employees, says Amrit Shahani, head of research at Coalition.

For the highly paid who remain, there is a growing income spread that mirrors the broader economy, says Babson College professor Tom Davenport. "The pay of the average managing director at Goldman will probably get even bigger, as there are fewer lower-level people to share the profits with," he says.

Complex trading algorithms, some with machine-learning capabilities, first executed trades where the price of what's being sold was easy to determine on the market, including the stocks previously traded by Goldman's 600 humans. Now complex things like currencies and credit, which are not traded on a stock exchange like the New York Stock Exchange but rather through less-transparent networks of traders, are coming in for more automation as well. To execute these trades, algorithms are being designed to emulate as closely as possible what a human trader would do, explains Coalition's Shahani.

Automation is hitting even the $700K-a-year investment bankers.

Goldman Sachs has already begun to automate currency trading and has found consistently that four traders can be replaced by one computer engineer, Chavez said at the Harvard conference. Some 9,000 people, about a third of Goldman's total staff, are computer engineers.

Next, Chavez said, will be the automation of investment banking tasks, work that traditionally has been focused on human skills like salesmanship and building relationships. Though those "rainmakers" won't be replaced entirely, Goldman has already mapped 146 distinct steps taken in any initial public offering of stock, and many are "begging to be automated," he said.

Reducing the number of investment bankers would achieve great cost savings for the firm. Investment bankers working on corporate mergers and acquisitions at large banks like Goldman make on average $700,000 a year, according to Coalition, and in a good year they can earn far more.

Chavez himself is an example of the rising role of technology at Goldman Sachs. His expertise in risk makes him suited to the task of CFO, a role typically held by accountants, Chavez told analysts on a recent Goldman Sachs earnings call.

"Everything we do is underpinned by math and a lot of software," he told the Harvard audience in January.

Goldman's new consumer lending platform, called Marcus, is aimed at consolidation of credit card balances and is run by software with no human intervention, according to Chavez. It was nurtured like a small startup within the firm and launched in just 12 months, he said.

It's a model Goldman is continuing, housing groups in "bubbles," some on the now-empty trading spaces in Goldman's New York City headquarters: "Those 600 traders, there is a lot of space where they used to sit," he said. *—Nanette Byrnes*

QUOTED

"When you get strangers bumping up against one another they behave more poorly."

— Raph Koster, a high-profile video-game designer, on why he expects trolling in virtual reality to become more common.

"It's possible to spend $8 million to piss off an entire city."

— Dean Rindy, a political consultant who helped sway voters in Austin, Texas, to enforce regulations on Uber and Lyft, which both left the city after the vote.

"For them, the clock is ticking, and there is a powerful sense of urgency."

— Mary Dunkle, vice president of educational initiatives at the National Organization for Rare Diseases, on the growing number of parents turning to gene therapy to cure their children.

BY THE NUMBERS

79 percent

Portion of U.S. homes that are suitable for rooftop solar power, according to a tool from Google called Project Sunroof.

$9,000

Cost to produce a pound of lab-grown chicken, according to Memphis Meats, one of a handful of startups trying to create meat without killing animals.

$100,000

Cost of a single xenotransplantation experiment on large animals, such as putting a pig's heart in a monkey.

1.1

Number of degrees Celsius by which 2016 global average temperatures exceeded preindustrial temperatures, according to the World Meteorological Organization.

Speeding Up Nuclear Research

An operator works in the control room at the MIT Nuclear Reactor Laboratory. Scientists there hope to accelerate advanced reactor research by building a prototype that piggybacks on their existing facility. Since the planned one-megawatt demonstration reactor would be incapable of sustaining a fission reaction on its own, the researchers believe they could avoid the lengthy regulatory approval process usually required for a standalone experimental prototype.

Photograph by Joshua Mathews

Upfront

The Next Big Encryption Fight May Not Happen in Public

Conflict over government access to encrypted data will inevitably reignite under President Trump.

How much power should the U.S. government have to compel technology companies to help it access their users' encrypted information? Last year's dramatic showdown between the FBI and Apple came to an end before the courts could give their answer, but the contentious debate is bound to flare up again before long in Washington. What might be in store this time around?

There is still a problem, according to law enforcement officials: the use of encryption is becoming more widespread, and products that do not allow even the providers themselves to access encrypted data, like Apple's iOS, unreasonably hinder investigations.

The government could again decide to go head to head in court with Apple and try to force the company to help investigators access an encrypted device, as it did with the iPhone of one of the San Bernardino shooters last year. If the FBI wants to avoid relitigating that dispute, it could instead target an encrypted messaging or e-mail service, applications that use a different form of encryption.

End-to-end encryption, an approach used by WhatsApp, iMessage, Signal, and others, prevents service providers from being able to read and hand over people's messages, as is possible with conventional e-mail and chat services. Much the way it tried to get Apple to build custom software that would help investigators access data stored on an iPhone, the FBI could try to make a messaging service provider provide some sort of technical assistance to help investigators read encrypted messages, says Andrew Crocker, staff attorney at the Electronic Frontier Foundation. The *New York Times* reported last March that the Justice Department was "privately debating" how to approach a standoff with WhatsApp over encryption.

The next big encryption-related legal fight will not necessarily play out in public, though, notes Crocker. The government could try to obtain a court order that prevents a company from speaking about a request by arguing that secrecy is necessary to stop a terrorist attack, for example. Pressuring Apple in public did not work out that well for the FBI because Apple was able to garner public support and the backing of privacy and civil liberties advocacy groups, including the Electronic Frontier Foundation.

More likely, the next high-profile encryption debate will happen in the new Congress. Near the end of Obama's presidency, the White House backed down from its position that law enforcement should get exceptional access to encrypted information, concluding that it wasn't possible without risking exposure to cybercriminals. Donald Trump, meanwhile, called for a boycott of Apple products when the company refused to help the FBI. His pick to head the Justice Department, Senator Jeff Sessions of Alabama, believes it is "critical that national security and criminal investigators be able to overcome encryption."

Broad proposals focused only on expanding government access to encrypted data will likely face strong opposition from lawmakers on both sides of the aisle, however. In December, a bipartisan "encryption working group" made up of influential members of the House of Representatives published a report concluding that Congress should not pass legislation that weakens encryption because such a law would work "against the national interest."

The report cited cryptographers and computer security experts who say it's not feasible to develop systems that give law enforcement exceptional access to encrypted information without also introducing undue risks to the security of all users of the technology. The report also pointed out that a U.S. law compromising encryption would be likely to cause users to shift to foreign-made products.

"There is no one-size-fits-all solution to the encryption challenge," wrote the report's authors, who include the chairs of the House Judiciary Committee and the Energy and Commerce Committee. They encouraged Congress to explore ways to help law enforcement take better advantage of the useful data that is already available, as well as tools like "legal hacking," and to foster more coöperation between law enforcement and technology companies. —*Mike Orcutt*

NISHANT CHOKSI

A Big Bet on Small Satellites

If all goes as planned, 2017 will set a new record for commercial launches of tiny spacecraft called CubeSats, each only a liter in volume and weighing less than two kilograms. The diminutive satellites have been used for over a decade in academic and government missions, but now investors and entrepreneurs are betting on new markets in imaging and telecommunications.

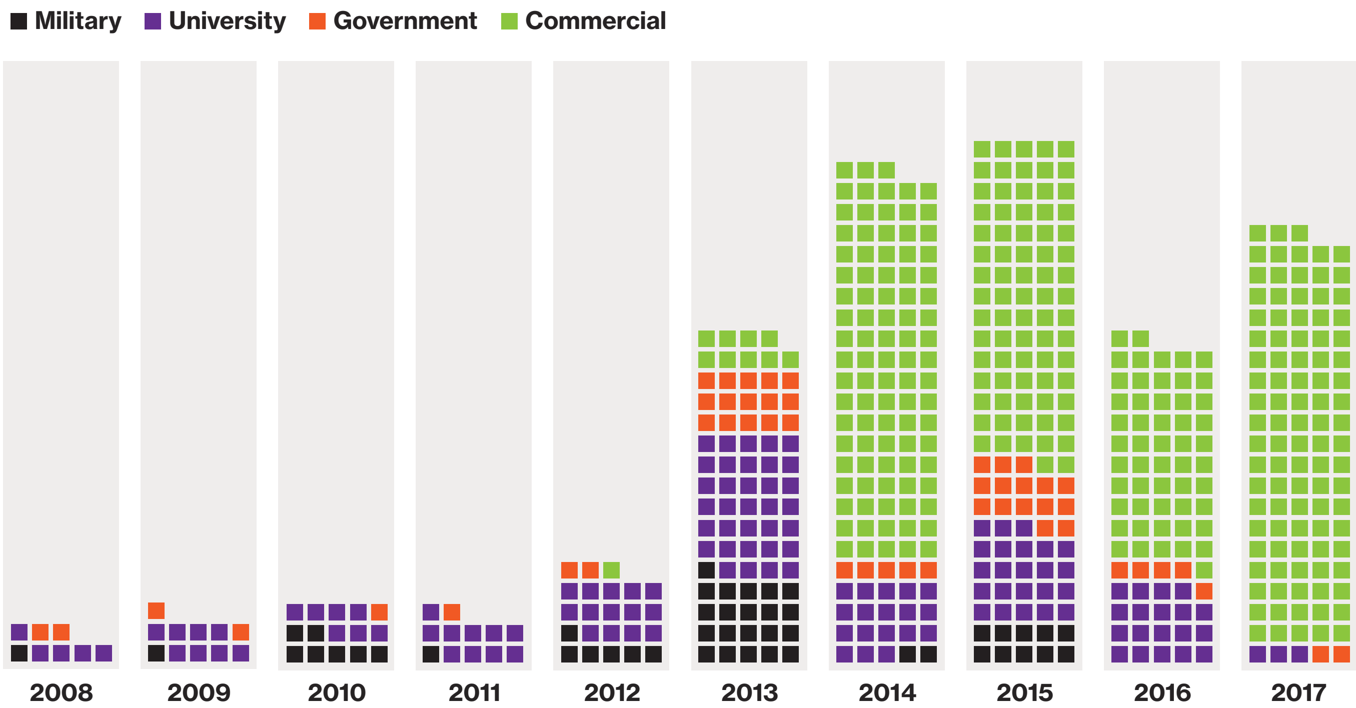

CubeSats Fly Low

CubeSats share low Earth orbit with roughly half the operational satellites in space. They can orbit for a few years before burning up. Much larger communications and imaging satellites orbit at 34,000 kilometers, and GPS satellites at 20,000 kilometers.

Rapid-Fire Data Capture

Industry leader Planet Labs says it can image roughly 50 million square kilometers per day using 61 orbiting CubeSats. Its newest models can capture up to 2.5 million square kilometers per day. By this summer, Planet Labs hopes to have the ability to capture the entire 150 million square kilometers of Earth's landmass every 24 hours, using 120 satellites.

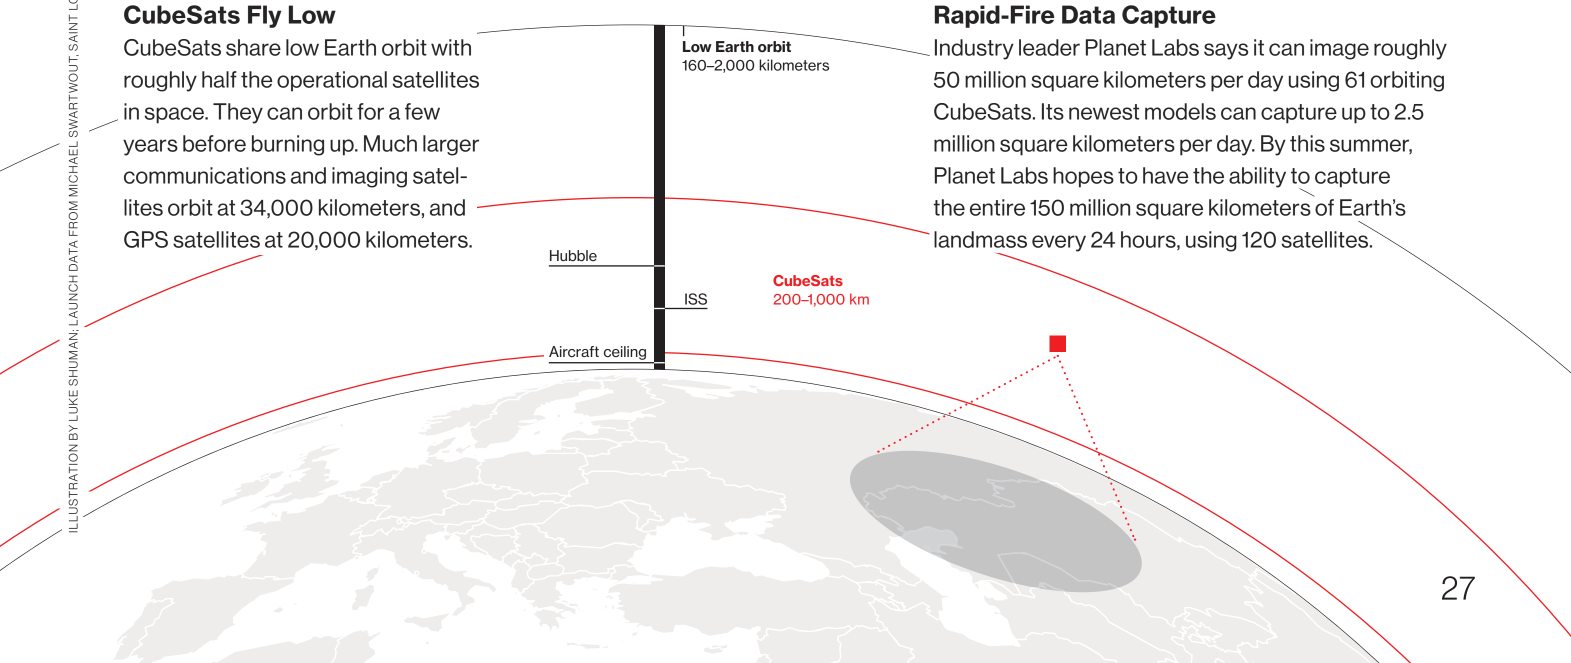

ILLUSTRATION BY LUKE SHUMAN; LAUNCH DATA FROM MICHAEL SWARTWOUT, SAINT LOUIS UNIVERSITY

The Growing Case for Geoengineering

As climate change accelerates, a handful of scientists are eager to move ahead with experiments testing ways to counteract warming artificially. Their reasoning: we just might get desperate enough to use this technology one day.

TATSURO KIUCHI

By James Temple

David Mitchell pulls into the parking lot of the Desert Research Institute, an environmental science outpost of the University of Nevada, perched in the dry red hills above Reno. The campus stares over the tops of the downtown casinos into the snow-buried Pine Nut Mountains. On this morning, wispy cirrus clouds draw long lines above the range.

Mitchell, a lanky, soft-spoken atmospheric physicist, believes these frigid clouds in the upper troposphere may offer one of our best fallback plans for combating climate change. The tiny ice crystals in cirrus clouds cast thermal radiation back against the surface of the earth, trapping heat like a blanket—or, more to the point, like carbon dioxide. But Mitchell, an associate research professor at the institute, thinks there might be a way to counteract the effects of these clouds.

It would work like this: Fleets of large drones would crisscross the upper latitudes of the globe during winter months, sprinkling the skies with tons of extremely fine dust-like materials every year. If Mitchell is right, this would produce larger ice crystals than normal, creating thinner cirrus clouds that dissipate faster. "That would allow more radiation into space, cooling the earth," Mitchell says. Done on a large enough scale, this "cloud seeding" could ease global temperatures by as much as 1.4 °C, more than the planet has warmed since the Industrial Revolution, according to a separate Yale study.

Big questions remain about whether it would really work, what damaging side effects might arise, and whether the world should risk deploying a tool that could alter the entire climate. Indeed, the suggestion that we should entrust the global thermostat to an armada of flying robots will strike many as preposterous. But the real question is: preposterous compared to what?

Without some kind of drastic action, climate change could be killing an estimated half-million people annually by the middle of this century, through famine, flooding, heat stress, and human conflict. Preventing temperatures from rising 2 °C above preindustrial levels, long considered the danger zone that should be avoided at all cost, now looks nearly impossible. It would mean cutting greenhouse-gas emissions by as much as 70 percent by 2050, and it may well require developing technologies that could suck megatons of carbon dioxide out of the atmosphere, according to the U.N.'s Intergovernmental Panel on Climate Change. But a growing body of research suggests that we probably will not have the time or technology to pull this off. Notably, even if every nation sticks to the commitments it's made under the politically ambitious Paris climate accords, global temperatures could still soar more than 5 °C by 2100.

"Everyone is looking at two degrees, but to me it's a pipe dream," says Daniel Schrag, director of the Harvard University Center for the Environment, who was one of President Obama's top advisors on climate change. "I fear we'll be lucky to escape four, and I want to make sure nobody ever sees six."

The difference between two and four degrees is another quarter-billion people without reliable access to water, more than a hundred million more exposed to flooding, and massive declines in worldwide crop yields, according to a study by the Committee on Climate Change, a London-based scientific group established to advise the U.K. government (see opposite page).

The idea that we could counteract these dangers by reëngineering the climate itself, techniques collectively known as geoengineering, began to emerge from the scientific fringes about a decade ago (see "The Geoengineering Gambit," January/February 2010). Now momentum behind the idea is building: increasingly grim climate projections have convinced a growing number of scientists it's time to start conducting experiments to find out what might work. In addition, an impressive list of institutions including Harvard University, the Carnegie Council, and the University of California, Los Angeles, have recently established research initiatives.

Few serious scientists would argue that we should begin deploying geoengineering anytime soon. But with time running out, it's imperative to explore any option that could pull the world back from the brink of catastrophe, says Jane Long, a former associate director at Lawrence Livermore National Laboratory. "I don't really know what the answer is," she says. "But I do believe we need to keep saying what the truth is, and the truth is, we might need it."

Dreams of dust

Mitchell works in a small, square office on the top floor of the Desert Research Institute. Stacks of scientific papers crowd his desk; journals and binders pack his bookshelf. Close-up images of delicate ice crystals hang from thumbtacks on the bulletin board above his computer monitor.

In the spring of 2005, during a sabbatical at the National Center for Atmospheric Research in Boulder, Colorado, Mitchell began exploring how the size of ice crystals affects cirrus clouds and the climate system. He and his colleagues found that bigger crystals, the type that tend to form in the presence of dust particles, produced fewer and thinner cirrus clouds.

That point stuck in Mitchell's brain. One morning shortly after returning to Nevada, he had a dream in which that insight morphed into a climate engineering scheme. He awoke wondering if deliberately adding dust in the areas where these clouds form would spawn these larger ice crystals, reducing cirrus coverage and releasing more heat into space.

2 °C TO 4 °C

Research shows it will be nearly impossible for the world to cut emissions fast enough to prevent 2 °C of warming above preindustrial levels. So what happens as global average temperature increases approach 4 °C? Hundreds of millions of additional people are exposed to climate dangers like flooding, hurricanes, drought, famine, heat stress, crop failure, and species loss. The mounting costs and fatalities could convince a growing number of nations to consider radical options like geoengineering.

POTENTIAL EFFECTS AS WARMING CLIMBS FROM 2 °C TO 4 °C

280 million more people lose access to adequate water

120 million more people exposed to major river floods

12 million more people subjected to coastal flooding

24% decline in global maize productivity

8% drop in spring wheat yields

34% of plant species lose half their suitable habitat

21% of mammal species lose half their habitat

HIGHER EMISSIONS ACCELERATE WARMING

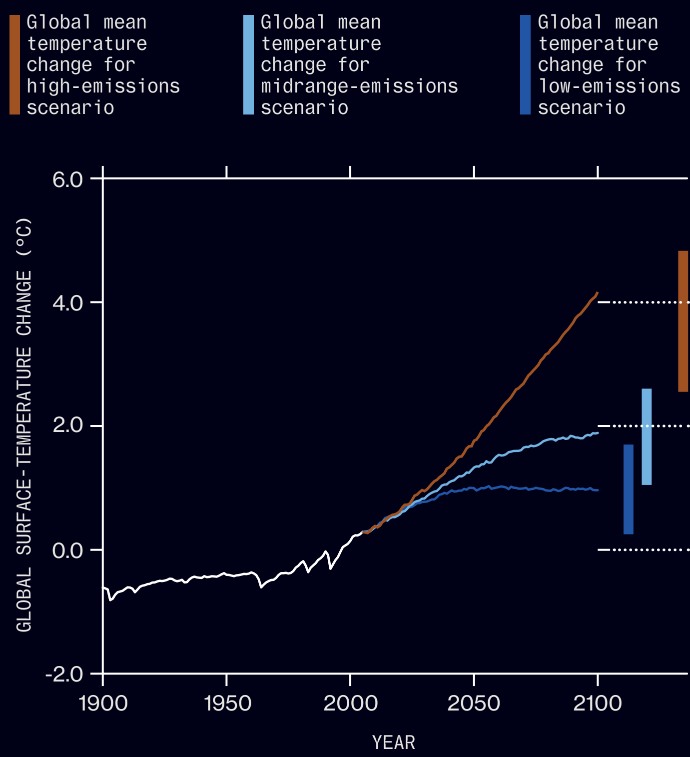

CHANGE IN AVERAGE PRECIPITATION (1986-2005 TO 2018-2100) UNDER LOW-EMISSIONS AND HIGH-EMISSIONS SCENARIOS

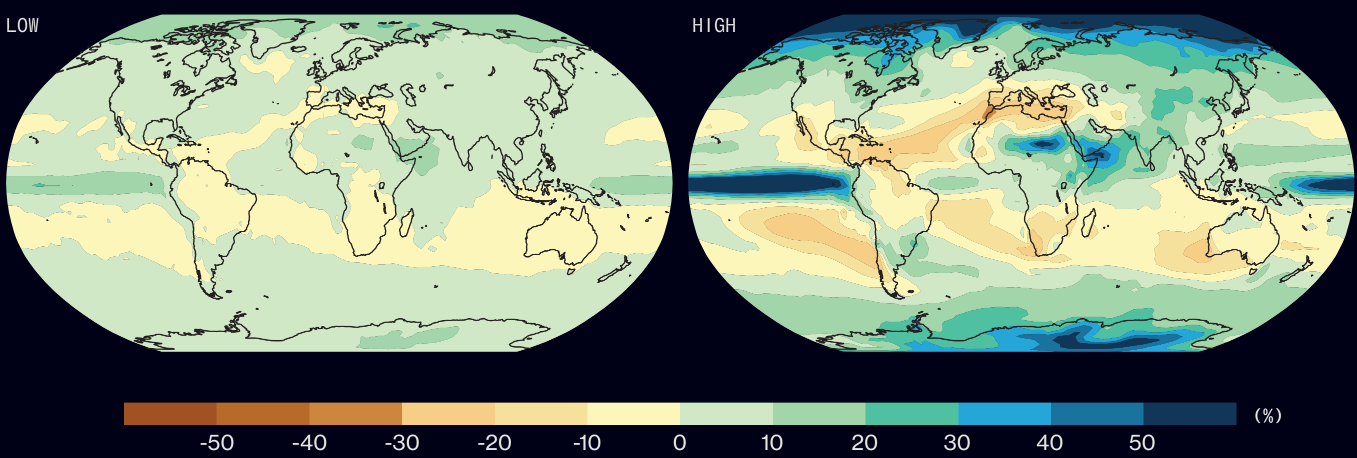

COMMITTEE ON CLIMATE CHANGE; U.N. INTERGOVERNMENTAL PANEL ON CLIMATE CHANGE

Though he had serious reservations about geoengineering, he decided to explore the idea. In 2009, he and a colleague published a paper suggesting that seeding cirrus clouds with tiny particles of bismuth tri-iodide, an inorganic compound that may break down into the necessary sub-micrometer size, might substantially offset climate change. More recently, Mitchell estimated that it would take around 160 tons of the material annually to seed clouds in the areas he has in mind, at a cost of about $6 million.

Not everyone agrees the proposal would work. A 2013 paper in *Science*, led by MIT atmospheric scientist Dan Cziczo, concluded that the formation of ice crystals around dust, known as heterogeneous ice nucleation, is already the dominant mechanism creating cirrus clouds. That might mean adding more dust would, on balance, create thicker clouds that trap more heat. The larger problem with the idea, Cziczo argues, is that clouds are the least understood part of the climate system. We do not have nearly enough knowledge about cloud microphysics, or accurate enough measurements, to precisely manipulate climate in this way, he says.

Without some kind of drastic action, climate change could be killing an estimated half-million people annually by the middle of this century.

But Mitchell's most recent research, relying on observations of ice crystal concentrations from NASA's Calipso satellite, has further convinced him that cloud seeding could work, as long as it's done in regions where cirrus clouds form primarily without dust particles. On the monitor in his office, Mitchell pulls up a page of maps from a paper he presented at the National Center for Atmospheric Research in late February. Navy- and light-blue dots, representing Cziczo's heterogeneous clouds, dominate the mid-latitudes, covering much of South America and Africa. But the higher latitudes are covered in red, yellow, orange, and green dots that indicate the sorts of clouds Mitchell has in mind.

The satellite images suggest that in very cold and humid conditions, toward the poles and particularly during winter, tiny ice crystals can form on their own, spontaneously, without dust. That suggests that cloud seeding could work, if it's targeted to those areas during those months. Mitchell even thinks he's come up with a way to get nature to carry out a field experiment to test his theory. During spring and winter, strong winds regularly stir up major dust storms in the deserts of Mongolia and the western edge of China. The fine particles blow across the Pacific and run into an atmospheric wave that rolls over the Rocky Mountains.

If Mitchell is correct, the dust should promote thinner cirrus clouds in an area where the thicker type otherwise tends to dominate. There was no way to properly observe this phenomenon—until late last year, when the National Oceanic and Atmospheric Administration launched a satellite equipped with some of the most powerful imaging technology ever launched into space, as well as sensors that can measure the temperatures of clouds. The satellite should be able to capture exactly what happens as the dust rides over the Rockies, detecting the subtle shifts under way in cloud microphysics.

Mitchell submitted a research proposal to NOAA last year, asking the agency to use the satellite to make such observations. He knows it's a long shot, particularly in light of the Trump administration's efforts to slash funding for climate science. But if NOAA agrees, the test could lend weight to his theory—or, of course, contradict it.

Another outdoor geoengineering experiment should occur even sooner.

By this time next year, Harvard professors David Keith and Frank Keutsch hope to launch a high-altitude balloon from a site in Tucson, Arizona. This will mark the beginning of a research project to explore the feasibility and risks of an approach known as solar radiation management. The basic idea is that spraying materials into the stratosphere could help reflect more heat back into space, mimicking a natural cooling phenomenon that occurs after volcanoes blast tens of millions of tons of sulfur dioxide into the sky (see "A Cheap and Easy Plan to Stop Global Warming," January/February 2013).

Scientists generally believe the technique would ease temperatures, but a lingering question is: what else will it do? Notably, volcanic eruptions have also significantly altered rainfall patterns in certain areas, and sulfur dioxide is known to deplete the protective ozone layer.

Keith has done extensive climate modeling to explore whether other materials, including alumina, diamond dust, and calcium carbonate, might have a neutral or even positive impact on ozone. During a conversation in his office at Harvard, he stressed that the experiments wouldn't constitute a test of geoengineering itself. But they would allow his group to subject its models to real-world data, revealing more about the relevant stratospheric physics and chemistry. "Theory alone doesn't tell you what will happen in the atmosphere," Keith says. "You can fool yourself if you don't go out and make direct measurements."

“The most likely scenarios for climate over longer time scales are devastating to future generations, absolutely devastating.”

Keith has already begun design work with the balloon company World View Enterprises, as well as discussions about the appropriate transparency and oversight for such outdoor experiments. The early flights would test the basic workings of the balloon, which would be tethered to a gondola equipped with propellers, sprayers, and sensors. But eventually the experiment would involve releasing a fine plume of materials, probably calcium carbonate, into the stratosphere. The balloon would then track that trail in reverse, allowing the sensors to measure how well the particles scatter sunlight, whether they coalesce or disperse, and how they interact with precursors to ozone.

Unknown unknowns

Full-scale geoengineering would inevitably involve some level of risk. We are likely to face a terrible choice between accepting the clear dangers of climate change and risking the unknowns of geoengineering. Alan Robock, a professor of environmental sciences at Rutgers, has published a list of 27 risks and concerns raised by the technology, including its potential to deplete the ozone layer and to decrease rainfall in Africa and Asia.

Ultimately, Robock worries that geoengineering may simply be too risky to ever try. “We don’t know what we don’t know,” he says. “Should we trust the only planet known to have intelligent life to this complicated technical system?” MIT’s Cziczo is blunter. “We know the problem is greenhouse gas, so the solution is you take the greenhouse gas out,” he says. “You don’t try to do something that we completely don’t understand.”

The reservations surrounding geoengineering research were on full display in late March as dozens of notable climate and social scientists gathered at the Carnegie Endowment for International Peace in Washington, D.C., for the Forum on U.S. Solar Geoengineering Research. Speakers highlighted a long list of unanswered, and perhaps unanswerable, questions about international governance: Who gets to decide when to pull the trigger? How do we determine “correct” average temperatures when the same ones will affect different nations in markedly different ways? Can one nation be held responsible for the negative effects of its geoengineering scheme on another country’s weather? Could these tools be used to deliberately attack a neighboring nation? And could conflicts over these questions tip into war?

“I have yet to hear any description of a future solar-geoengineered world that sounds to me anything other than dystopian or highly unrealistic,” said Rose Cairns, a research fellow at the University of Sussex, who joined the morning discussion from England by Skype.

But Harvard’s Schrag argued the opposite: that the scariest version of the future may be one where geoengineering is never developed or deployed. “I don’t think people understand just what we’re up against with climate,” he said. “The most likely scenarios for climate over longer time scales are devastating to future generations, absolutely devastating.”

As he flashed slides highlighting the dramatic loss of sea ice in the Arctic and Antarctic in recent months, Schrag stressed that climate change is already causing visible impacts faster than anyone expected. He added that it’s difficult to foresee any scenario where we can cut greenhouse-gas levels fast enough to avoid far worse dangers: the amount we’ve already released is likely to lock in another degree of warming even if we halt emissions tomorrow, he said.

To his mind, these hard realities mean we need to try to answer the difficult questions geoengineering poses. “It is still, in every case that I’ve seen, better than the alternative of just letting the climate warm,” he said. “Given the trajectory of the world, and the difficulty of reducing emissions, this is something we really need to understand.”

The power of fear

Mitchell was opposed to geoengineering for most of his career. The idea that humankind should tinker with the finely tuned climate system struck him as impossibly arrogant. But like other researchers who spent decades staring at increasingly frightening projections while the world ignored the loudest warnings scientists knew how to sound, he reluctantly changed his view.

It could take decades to learn which geoengineering methods might work, whether environmental side effects can be minimized, and whether it’s ultimately too dangerous to try. The longer we wait to begin serious research, the greater the risk we’ll deploy an unsafe tool in the face of sudden climate shocks, or not have one in hand when we need it. And no one really knows when that might be.

Says Mitchell, “The need for climate engineering could be coming faster than we realize.”

James Temple is MIT Technology Review*’s senior editor for energy.*

TREATING ADDICTION WITH AN APP

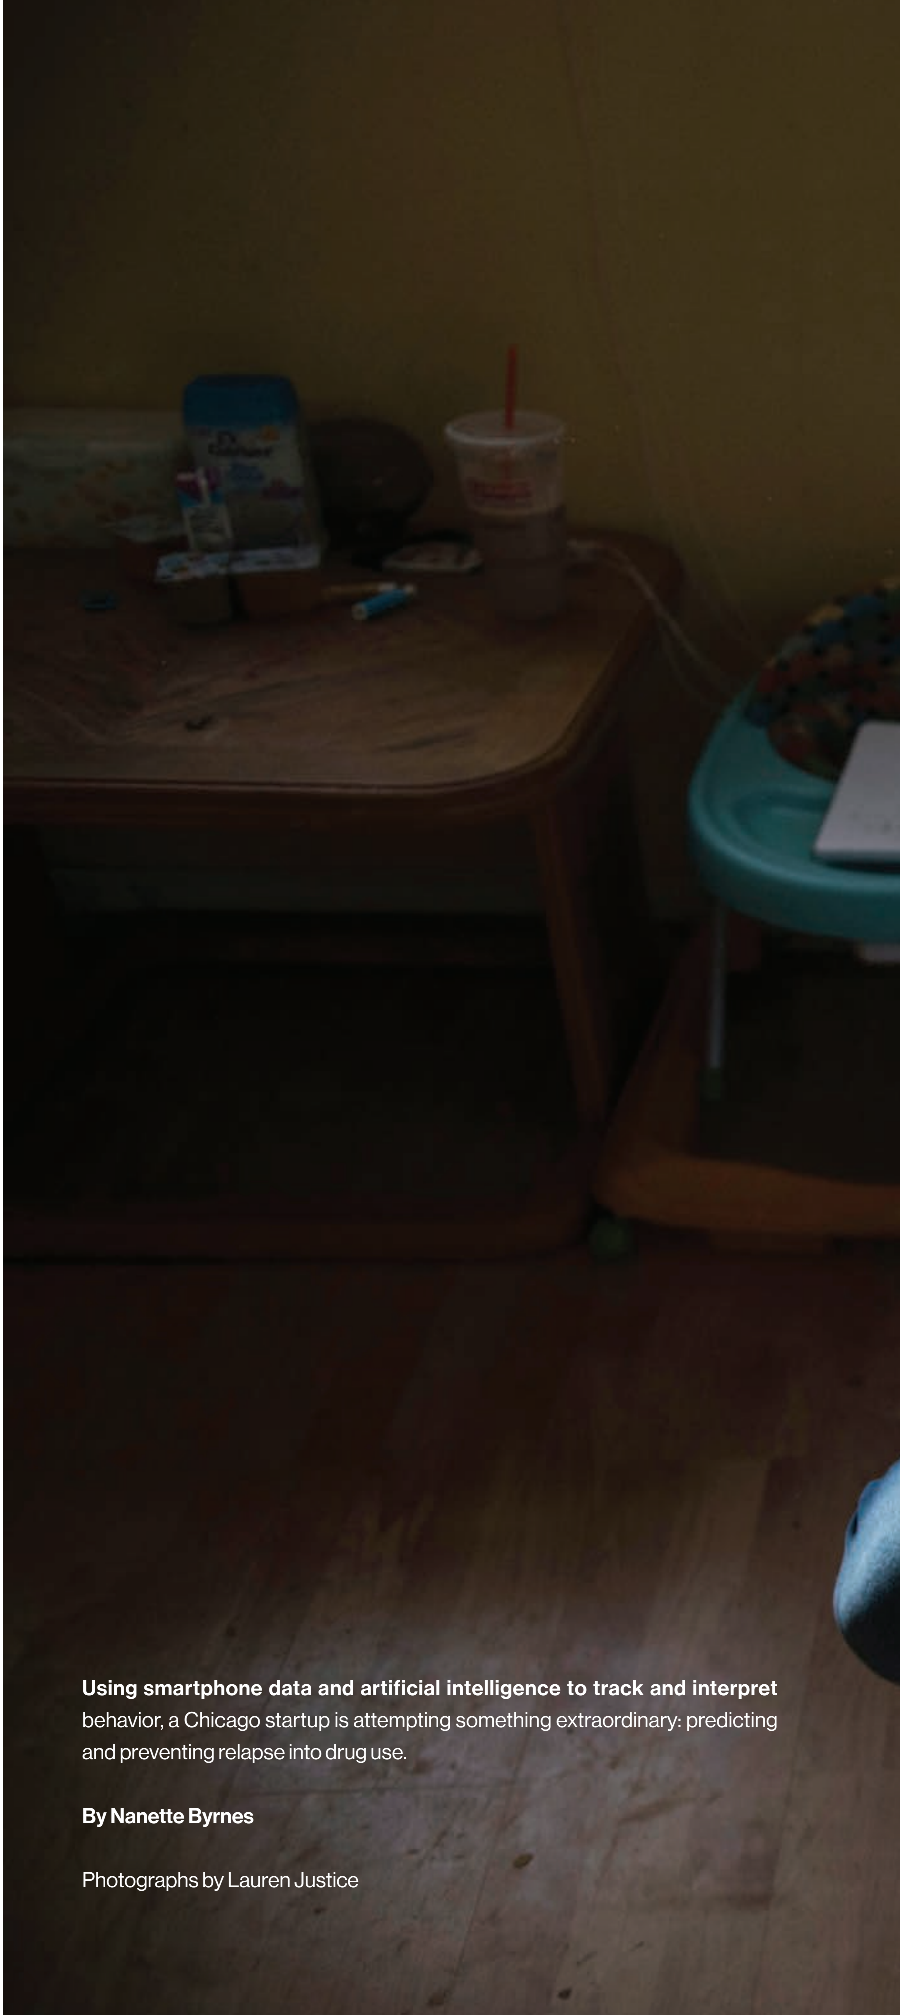

Using smartphone data and artificial intelligence to track and interpret behavior, a Chicago startup is attempting something extraordinary: predicting and preventing relapse into drug use.

By Nanette Byrnes

Photographs by Lauren Justice

Tasha Hedstrom has been using the Triggr platform to help her manage her recovery from opioid addiction.

WHEN I SPOKE TO TASHA HEDSTROM THIS WINTER, SHE HAD BEEN SOBER FOR MORE THAN 61 DAYS. AFTER STRUGGLING WITH OPIOID ADDICTION FOR 15 YEARS, HEDSTROM IS TAKING VIVITROL, A DRUG THAT BLOCKS THE PLEASURABLE EFFECTS OF OPIOIDS AND REDUCES CRAVINGS. SHE GOES TO A COURT-MANDATED RECOVERY PROGRAM THREE DAYS A WEEK AND TRACKS HER PROGRESS ON A PHONE APP SHE FOUND ON FACEBOOK, CALLED TRIGGR HEALTH.

H E D S T R O M

S A Y S

she has never found peer support programs like Narcotics Anonymous helpful. “I don’t like the atmosphere. I feel like people are talking about using and glorifying it,” she says. “I don’t like telling my story a million different times.”

Triggr has been a different way to access support. In addition to tracking the number of days she has been in recovery, the app connects Hedstrom to a team of recovery coaches, who chat with her periodically throughout the day by text and app message. If she has not contacted Triggr for a full day, the team contacts her. Generally, they talk about how her day is going or goals she has set for herself, but recently they helped her through an unexpected challenge. A stranger followed her car into a lot and parked next to her, then offered her drugs. Not sure what to do, Hedstrom texted Triggr. “It’s not just about addiction,” she says. “It’s like we’re on a friend basis. You need to have backup supports.”

In 2015, 33,000 people in the United States died from opioid overdoses—the highest number ever recorded, and more than double the 2005 figure, according to the National Institute on Drug Abuse. More than half a million hospitalizations related to opioid dependence occur each year, at a cost of $15 billion, according to a recent study. Tens of billions more are spent on clinics and other treatment.

In total, 23 million Americans have a substance use disorder involving illicit drugs or alcohol, according to 2013 data collected by the Substance Abuse and Mental Health Services Administration, part of the U.S. Department of Health and Human Services. But fewer than 20 percent of those who need treatment will receive it. And while the most common form of treatment, Alcoholics Anonymous or Narcotics Anonymous, can be quite effective for some, according to one survey 75 percent of people in such programs relapse in their first year. Though a wide range of treatments are available, says James R. McKay, an expert in addiction and a professor of psychology at the University of Pennsylvania, “there are people for whom none of those things really work.”

New technologies are offering yet another option, making use of the computers we carry in our pockets. Of the 165,000 smartphone apps for health care, mental-health apps are the single largest category, including hundreds of addiction-related options offering inspirational quotes, directions to nearby AA meetings, hypnosis guides, and online peer support groups.

Triggr is more ambitious. Using data collected from smartphones, the company aims not only to help people handle cravings and the stresses that trigger drug use but to actually predict when someone is going to relapse and intervene. Triggr collects clues from things like screen engagement, texting patterns, phone logs, sleep history, and location. Those are combined with information gathered from participants’ communications with the startup’s staff on its platform—such as drug preference, drug history, and the presence of dangerous words like “craving” or “stress”—and fed into a series of algorithms. The system has access to general information about other texting and e-mail activity but not to the content of private texts or calls. Using machine learning, it searches for patterns that point to an increased likelihood of relapse. When the likelihood rises to a dangerous level, a member of the recovery team steps in or alerts a customer’s outside care team.

Neither the company nor clients will say how much the platform costs to use, though in some pilot projects Triggr seems to be charging very little or nothing. Hedstrom downloaded the app for free but now pays a monthly fee for using the system, which she says is less than two dollars a day. The most promising way for Triggr to make money could be to share in the financial savings use of the app could offer the insurers and government agencies that pay the medical costs associated with addiction. An initial 30-day inpatient treatment can cost $17,000, and emergency room visits and other associated costs add up quickly.

Chris Olsen, a partner at the venture capital firm Drive Capital, one of Triggr’s investors, says it has been estimated that Ohio Medicaid is spending as much as $5 billion a year on hepatitis C infections, which are strongly correlated with injection drug use. “If we can reduce that,” says Olsen, “I just believe there will be a revenue model down the line.” Among those using the app today are patients who have been through rehab at Sprout Health Group, a chain of addiction treatment centers headquartered in New Jersey. Sprout CEO Arel Meister-Aldama says that before Triggr, a patient who had been in a full-time program for 45 days on average and then returned to the community would have been tracked with periodic phone calls and invited to alumni events, but it was hard to know how people were really doing. Now Sprout’s counselors get alerts from Triggr when a patient seems at risk. “There are false alarms, but often we’ll catch people on the way to their drug dealers. Or they’ll be sitting outside a bar thinking about going in,” says Meister-Aldama.

Sprout’s readmission rates have actually gone up since the company started using Triggr, but overall cost per patient has declined. That’s because its counselors have been able to help patients earlier, avoiding expensive stays in treatment facili-

ties and emergency treatment. With the data he is getting from Triggr, Meister-Aldama says, he has a better understanding of what it will cost to treat each patient. He expects that in the future he will be able to agree to flat payments per patient instead of charging fees based on services.

The platform Meister-Aldama has found so useful wouldn't work without ubiquitous smartphones and recent advances in machine learning. And it wouldn't exist at all if it had not been for one college student's pain—and her mother's timely intervention.

MOTHERLY INTUITION

Art by Hedstrom's sons Riley, four years old, and Jeremiah, 11 months, decorates the walls of her home.

Having struggled with addiction for 15 years, she has found that 12-step groups are not her thing. Triggr, though, has become a valued support.

John Haskell, Triggr's cofounder and CEO, came up with the idea for the app and the broader system of care during a challenging period in his own life. While an undergraduate at Stanford, he battled manic depression, spending five years at school without earning a degree. And one of his friends at Stanford struggled with mental-health problems and substance abuse. She got to a point at which she did not want to continue with treatment and considered suicide. At a particularly critical moment, her mother called. The call set her daughter on a more positive path, and when Haskell asked the mother what had prompted the call just at that moment, she attributed it to "motherly intuition."

Motherly intuition was something Haskell thought should be reproducible with the help of smart technology.

"She knew something was wrong. She could feel it. But what was particularly interesting about that experience was that it was all these data points. And all trackable on your phone," he says. For example, his friend had always loved Words with Friends, an online multiplayer game similar to Scrabble, but she had stopped playing. She was sending texts in the middle of the night, an obvious sign she was not sleeping. "The concept of intuition is purely a data question," says Haskell. "Why can't you scale motherly intuition?"

Six years later, Haskell's motherly-intuition machine occupies two long white tables in a second-floor walk-up in Chicago's River North neighborhood. At one table sit a small group of programmers and data scientists, many with backgrounds at larger companies including Google, building the app and its platform. On the other side of a partial partition wall, at an identical white

table, sit the recovery group, a team of four to five people who interact with participants on the platform. Everyone faces a computer screen.

The technologists work with the data that sensors are pulling off participants' phones as well as from their interactions with the recovery team, identifying patterns that signal a move in the wrong direction. Twenty-four hours a day, seven days a week, Triggr actively watches over everyone on the platform, with a single member of the recovery team following 500 people at any time. Each participant has a rating on a scale of 1 to 10 based on the patterns Triggr's algorithm is tracking. A 1 means things are going very well. A 10 is an alert that the person is exhibiting a pattern of behavior that may be on the edge of relapse and needs to be contacted.

Most staff communication with clients takes place via text or app messaging. Without the clues they might get in person from eye contact and body language, or on the phone from someone's tone of voice, the team relies heavily on alerts from its algorithms.

ARE YOU A ROBOT?

Triggr's machine-learning systems have made the platform smarter over time by studying both those interactions with participants and the millions of data points collected from their smartphones. The systems search for anomalies, breaks from a client's typical routine. As more people use the system and more data is gathered and studied, the ability to see signs of a potential relapse improves. Eighty-five percent accurate a year ago, Triggr can now predict with 92 percent accuracy when a client is likely to slip in the next three days. The early intervention such predictions make possible is significantly improving clients' results, the company says.

The messiness of the data is what convinced Triggr's data scientist, John Santerre, that machine learning could be effective against the problem. Some of the most important warning signs of an impending slip have nothing directly to do with drugs or alcohol. Instead, they're life events, like the death of a family member or another user, an affair, an issue with housing. Just one deviation from a client's normal routine—something

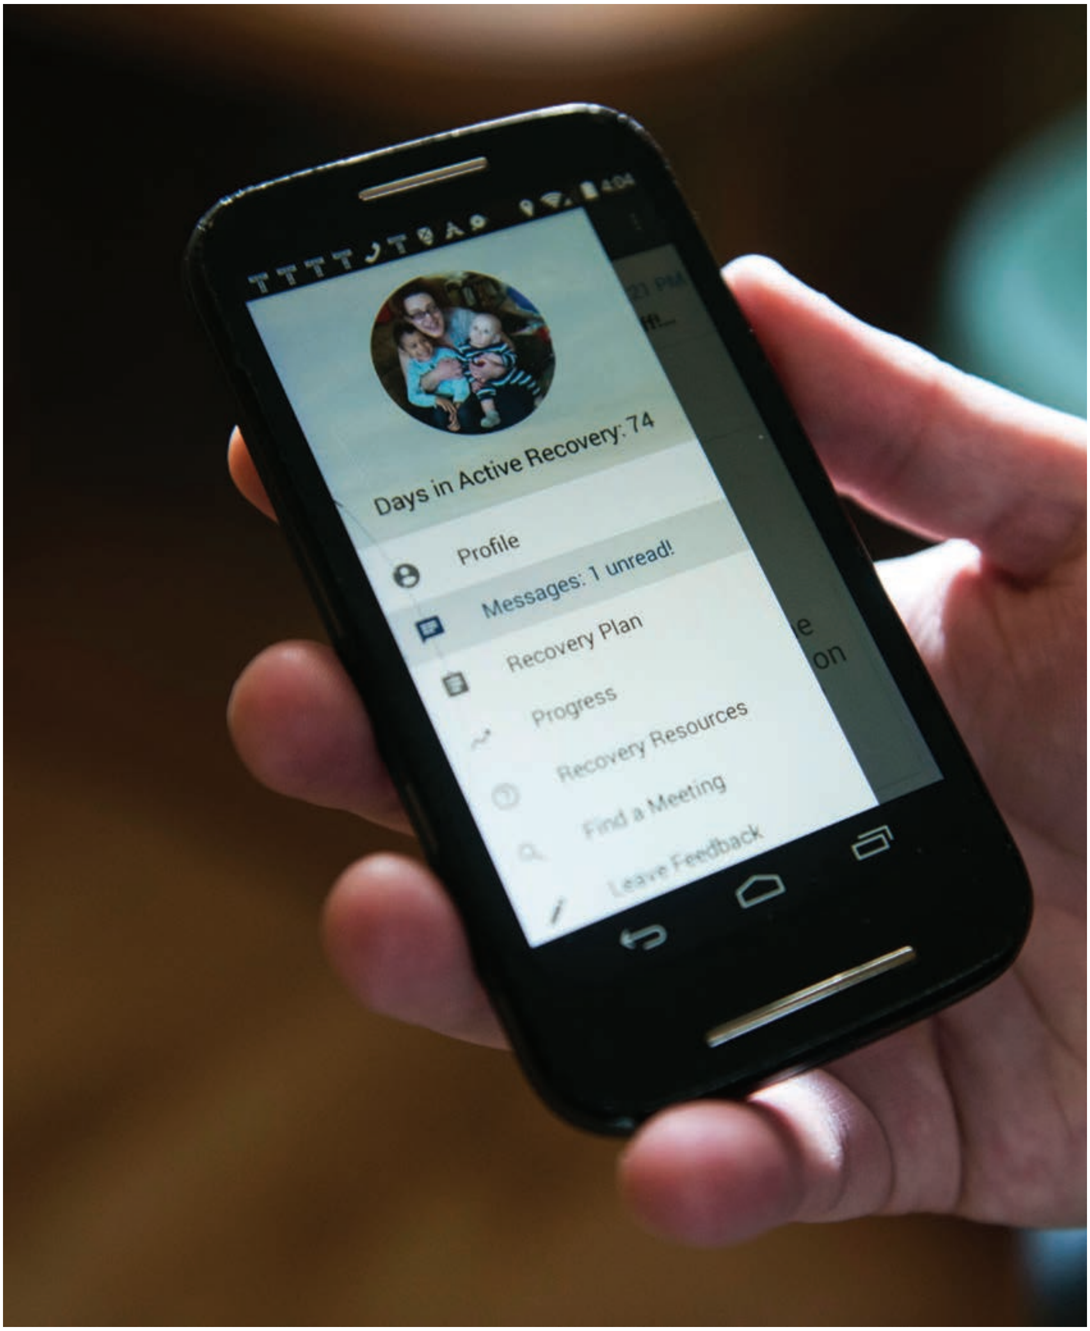

Hedstrom doesn't check her Triggr app at specific times of day, but when she does check, she says, she gets a response from the recovery staff right away.

as small as a text that comes in at an unusual time—increases the chance of relapse in the next few days. Triggr does not even need to know whom that text comes from or what it says. The interruption of routine is the critical clue.

Triggr is collecting every piece of data it can on how to help people resist an urge as it swells and then drops off, and has taken on the tricky task of building a system designed to work with minimal human input while producing a service customized to each participant. While algorithms may determine that a slip is coming, intervening to stop it isn't necessarily suited to automation. "Our goal is to make it as human as possible," says Haskell. Still, clients do sometimes ask the recovery coaches if they are robots. Tasha Hedstrom did; Triggr responded by asking if *she* was a robot. Humor is one of the techniques the algorithm has determined work well with some participants.

The coaches are always testing messages sent to clients in response to different types of issues. Those that resonate are shared with the engineering team; when a similar call comes in later, the system will know to suggest that effective response. Once Triggr determines that a person is in danger of relapsing, it's time for the really hard part: intervention to stop the self-destructive behavior. Humans do oversee the interaction, but when someone's risk is rising, a member of the recovery team is automatically alerted to the most effective way to reach out to that client and the type of message to which he or she is most likely to respond. This is as close to Haskell's idea of digital intuition as Triggr has come so far.

NO SMARTPHONE

A big focus for Haskell is developing connections to community service organizations, and on a wet morning in January, he was standing in a conference room in Framingham, Massachusetts, excitedly explaining the app to a group of counselors from the South Middlesex Opportunity Council (SMOC), a local nonprofit. SMOC had just launched Triggr as part of a program to connect with drug users in the emergency room after they have overdosed. Like many parts of the Northeast, the Midwest, and Appalachia, Framingham is suffering a rising number of opioid-related overdoses: they now average 10 a month.

Some counselors in the room worried that not all potential clients have smartphones. Others wanted a service Triggr does not offer: alerts when a client has contacted a drug dealer or used drugs again. Haskell had answers for every question, but a month and a half after the presentation, Krystin Fraser, who is running the grant, said that of the first eight people who signed up, only one agreed to download Triggr. Some do not have a smartphone, she explained, while others simply do not want someone watching them. Over the next month, 13 more people signed up for the app.

Most health apps are not regulated by the Food and Drug Administration and the company has chosen not to publish any clinical trials of its platform, something it is not required to do. It is tracking the long-term outcomes for people who use Triggr, and its decision does put the burden on the company to show that it really has made something extraordinary. It is in a crowded field. "There's been a glut of mental-health apps, most of questionable use and efficacy," says John Torous, a director of the digital psychiatry program at Boston's Beth Israel Deaconess Medical Center. Torous is part of a study using passive phone data to follow people suffering from schizophrenia, a mental disorder that is quite different from addiction but can feature similar underlying behavior, such as disrupted sleep. "People underestimate how complex it is to work with this data," says Torous. "We've had mass-market smartphones for 10 years and we still haven't revolutionized mental-health care. If this were as easy as building an app, in 10 years it would have been done. People are complex. We can collect all this data, but how do we analyze it in a valid way?"

Jukka-Pekka Onnela, a professor of biostatistics at Harvard's T.H. Chan School of Public Health and Torous's collaborator on the schizophrenia study, is more optimistic. As people use phones for more and more daily needs such as schedules, navigation, and communication, the data from these devices becomes "very, very powerful," Onnela says. That's especially so for conditions where behavior is strongly influenced by a person's surroundings and recent history, as it is for people with psychological disorders or addiction.

When they're awake, people's phone screens may be on more than 20 times an hour. Onnela has found that frequency to be a reliable indicator of sleep patterns, something essential to understanding psychological illness and treating it.

"In the past a lot of measurement has been confined to labs or doctor's offices," says Onnela. "What we are trying to do is to capture symptoms in the wild, the way people actually experience their lives."

Nanette Byrnes is MIT Technology Review*'s senior editor for business.*

One of Desktop Metal's 3-D printers makes a part out of a steel alloy, demonstrating its ability to churn out complex structures.

THE 3-D PRINTER THAT COULD FINALLY CHANGE MANUFACTURING

DESKTOP____ME
ITS__MACHINES_
DESIGNERS_AND_
ERS___A___PRAC
A F F O R D
PRINT____METAL_

It's less than two months before his company's initial product launch, and CEO Ric Fulop is excitedly showing off rows of stripped-down 3-D printers, several bulky microwave furnaces, and assorted small metal objects on a table for display. Behind a closed door, a team of industrial designers sit around a shared work desk, each facing a large screen. The wall behind them is papered with various possible looks for the startup's ambitious products: 3-D printers that can fabricate metal parts cheaply and quickly enough to make the technology practical for widespread use in product design and manufacturing.

TAL___THINKS _WILL___GIVE _MANUFACTUR- TICAL___AND ABLE__WAY__TO _PARTS.

By David Rotman

PHOTOGRAPHS BY GRANT CORNETT

The company, Desktop Metal, has raised nearly $100 million from leading venture capital firms and the venture units of such companies as General Electric, BMW, and Alphabet. The founders include four prominent MIT professors, including the head of the school's department of materials science and Emanuel Sachs, who filed one of the original patents on 3-D printing in 1989. Still, despite all the money and expertise, there's no guarantee the company will succeed in its goal of reinventing how we make metal parts—and thus transforming much of manufacturing.

As Fulop moves about the large, open workspace, his excitement and enthusiasm seem tempered by anxiety. The final commercial printers are not yet ready. Employees are busy tinkering with the machines, and fabricated test objects are scattered about. Progress is being made, but it's also obvious that the clock is ticking. In a corner near the front door and entrance area, the floor is empty and taped off; soon the space needs to be filled with a mockup of the company's planned booth for an upcoming trade show.

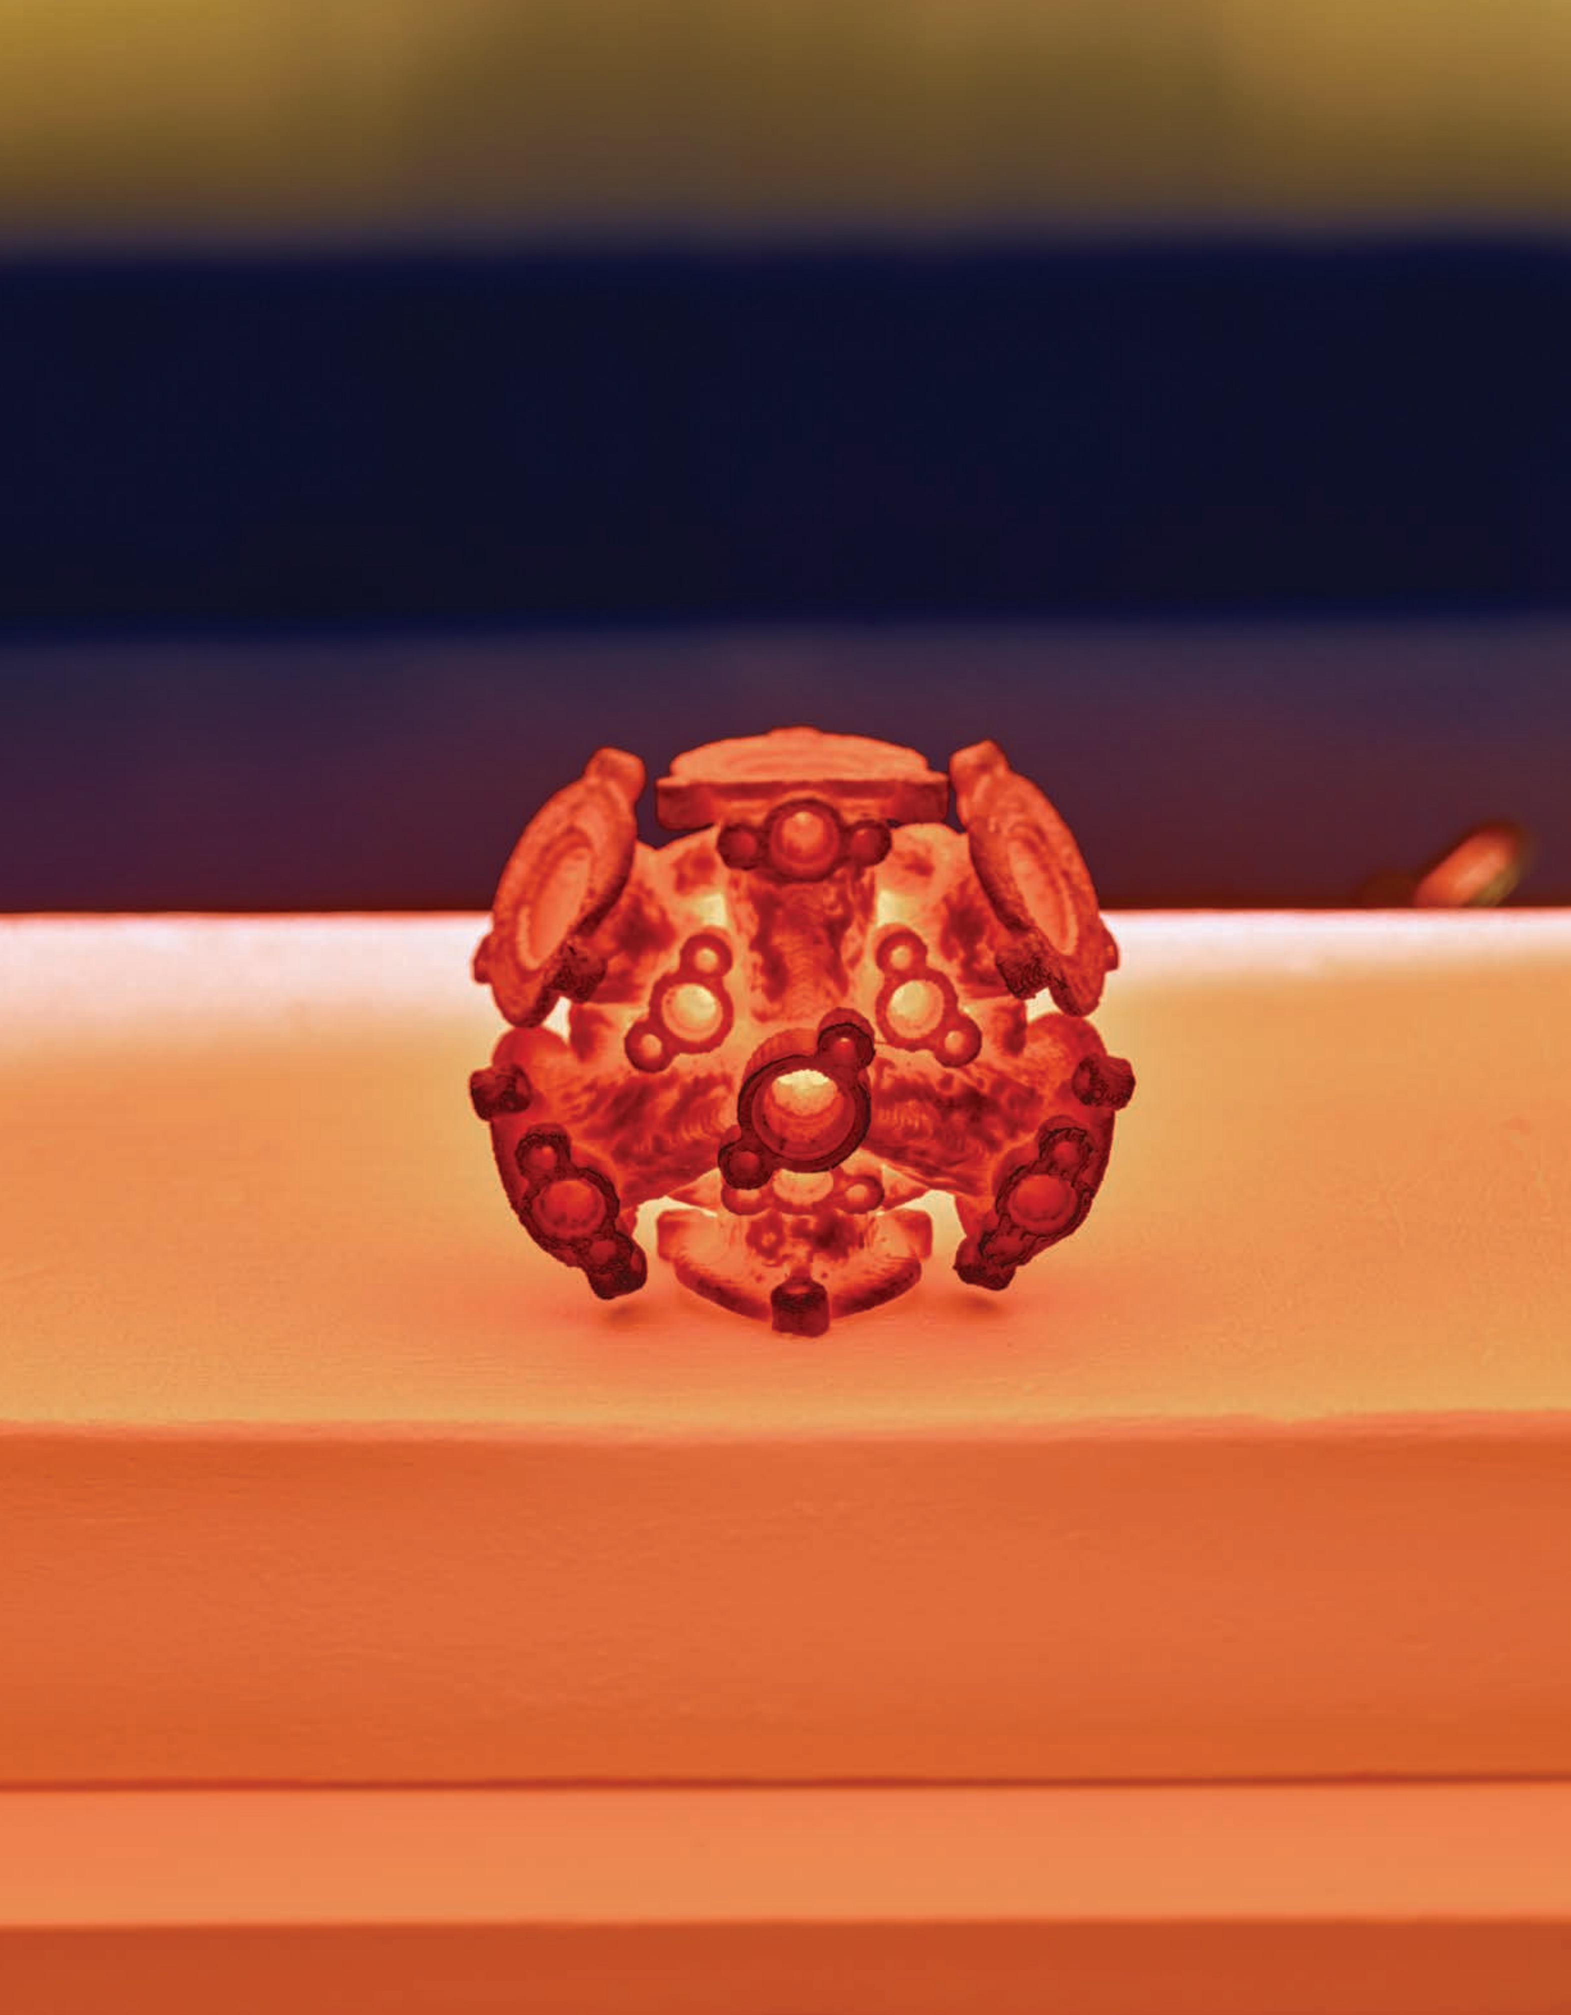

If it succeeds, Desktop Metal will help solve a daunting challenge that has eluded developers of 3-D printing for more than three decades, severely limiting the technology's impact. Indeed, despite considerable fanfare and evangelical enthusiasts, 3-D printing has, in many ways, been a disappointment.

Hobbyists and self-proclaimed makers can use relatively inexpensive 3-D printers to make wonderfully complex and ingenious shapes out of plastics. And some designers and engineers have found those machines useful in mocking up potential products, but printing polymer parts has found little use on the production floor in anything but a few specialized products, such as customized hearing aids and dental implants.

Though it is possible to 3-D-print metals, doing so is difficult and pricey. Advanced manufacturing companies such as GE are using very expensive machines with specialized high-power lasers to make a few high-value parts (see "Additive Manufacturing" in our 10 Breakthrough Technologies list of 2013). But printing metals is limited to companies with millions to spend on the equipment, facilities to power the lasers, and highly trained technicians to run it all. And there is still no readily available option for those who want to print various iterations of a metal part during the process of product design and development.

The shortcomings of 3-D printing mean the vision that has long excited its advocates remains elusive. They would like to create a digital design, print out prototypes that they could test and refine, and then use the digital file of the optimized version to create a commercial product or part out of the same material whenever they hit "make" on a 3-D printer. Having an affordable and fast way to print metal parts would be an important step in making this vision a reality.

It would give designers more freedom, allowing them to create and test parts and devices with complex shapes that can't be made easily with any other production method—say, an intricate aluminum lattice or a metal object with internal cavities. It could eventually enable engineers and materials scientists to create parts with new functions and properties by depositing various combinations of materials—for example, printing out a magnetic metal next to a nonmagnetic one. Beyond that, it would redefine the economics of mass production, because the cost of printing something would be the same regardless of how many items were produced. That would change how manufacturers think about the size of factories, the need for backup inventory (why keep many parts in stock if you can simply and quickly print one out?), and the process of tailoring manufacturing to specialized products.

This is why there has been a race to turn 3-D printing into a new way to produce parts. Longtime suppliers of 3-D printers, including Stratasys and 3D Systems, are introducing increasingly advanced machines that are fast enough for manufacturers to use. Last year, HP introduced a line of 3-D printers that the company says will allow manufacturers to prototype and make products with nylon, a widely used thermoplastic. And last fall, GE spent over a billion dollars on a pair of European companies specializing in 3-D-printing of metal parts.

But the real competition for Desktop Metal is probably not from the growing number of companies in 3-D printing. For one thing, the 3-D printers from HP, Stratasys (an investor in Desktop Metal), and 3D Systems mainly use various types of plastics, not the range of metals Fulop's company wants to use in its printers. And GE's high-end machines overlap little with Desktop Metal's market ambitions. Instead, the real competitors for Desktop Metal are more likely to be established metal-processing technologies. Those include automated machining techniques—such as the method used to make the ultra-thin aluminum back casing of iPhones—and a rapidly growing practice called metal injection molding, a common way to mass-produce metal products.

In other words, rather than merely trying to outdo other 3-D printers, Desktop Metal will have the tough task of converting manufacturers away from production methods that are at the heart of their businesses. But the very existence of this large, established market is what makes the prospect so intriguing. Making metal parts, says Fulop, "is a trillion-dollar industry." And even if 3-D printing wins only a small portion of it, he adds, it could still represent a multibillion-dollar opportunity.

Too hot to print

Look around. Metals are everywhere. But whereas 3-D printing has been widely used in making plastics, the technology's use in making metal parts "has been narrowly confined," says Chris Schuh, head of materials science and engineering at MIT and cofounder of Desktop Metal. "Metal processing is more of an art. It's a very challenging space."

Making metal objects using 3-D printing is difficult for several reasons. Most obvious is the high temperature required for processing metals. The most common way to print plastics involves heating polymers and squirting the material out the printer noz-

A hydraulic manifold is processed inside a microwave furnace, which uses temperatures up to 1,400 °C to sinter the steel part. Such a part is too complex to make with conventional methods.

This steel propeller has just been printed. Between the propeller's blades and the metal support is a thin line of ceramic, which will turn to sand during the sintering process, allowing the finished part to be easily separated from the support.

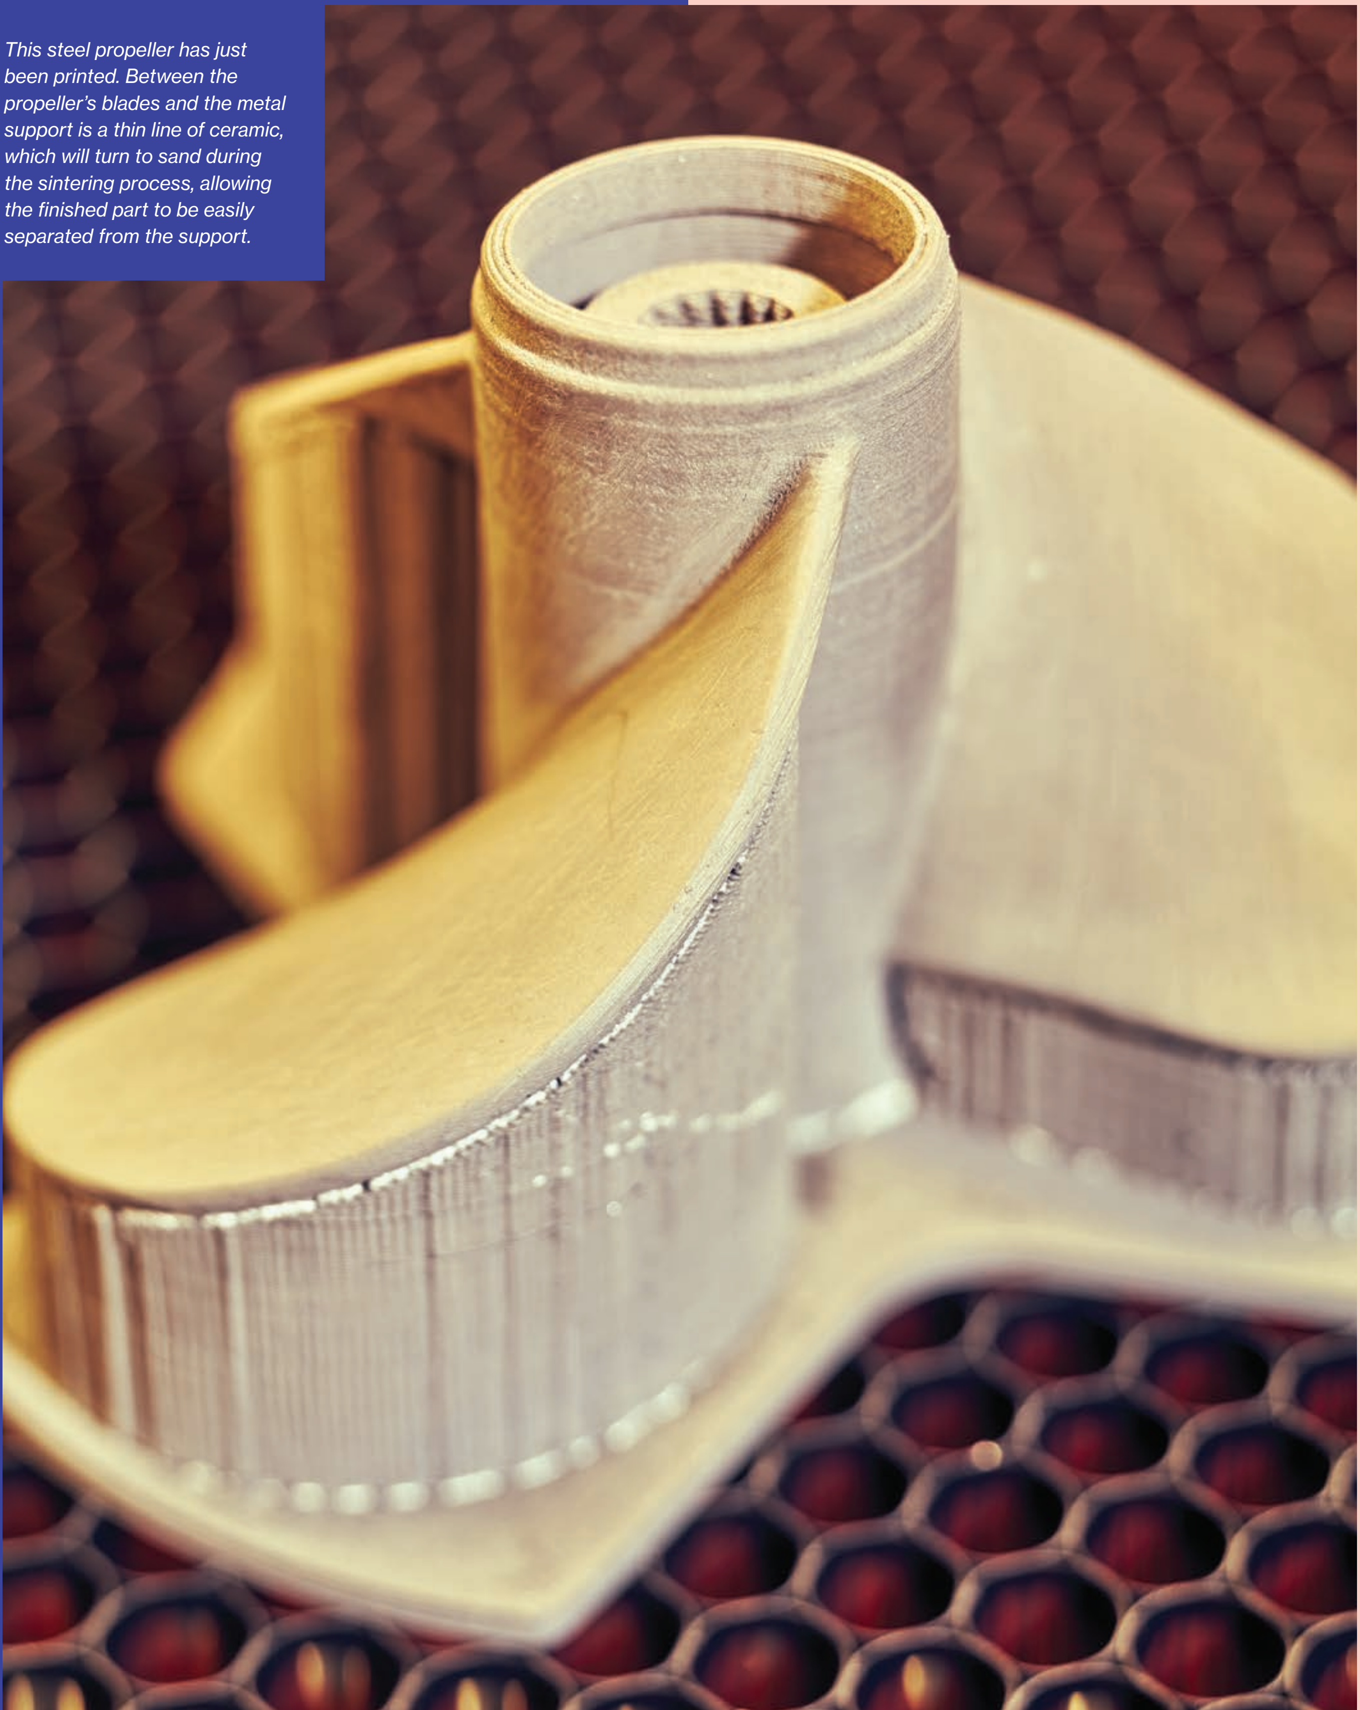

The propeller after processing provides an example of a high-performance part that can be made with 3-D printing. Engineers can use the method to prototype and optimize different designs.

zle; the plastic then quickly hardens into the desired shape. The process is simple enough to be used in 3-D printers that sell for around $1,000. But building a 3-D printer that directly extrudes metals is not practical, given that aluminum melts at 660 °C, high-carbon steel at 1,370 °C, and titanium at 1,668 °C. Metal parts also have to go through several high-temperature processes to ensure the expected strength and other mechanical properties.

To make a 3-D printer fast enough to be used in manufacturing metal objects, Desktop Metal turned to a technology that dates back to the late 1980s. That's when a team of MIT engineers led by company cofounder Sachs filed a patent for "three-dimensional printing techniques." It described a process of putting down a thin layer of metal powder and then using ink-jet printing to deposit a liquid that selectively binds the powder together. The process, which is repeated for hundreds or thousands of layers to define a metal part, can make ones with nearly unlimited geometric complexity. In the most common application of the technology, the binder acts like a glue. However, it can also be used to locally deposit different materials in different locations.

The MIT researchers knew their printing method could be used to make metal and ceramic parts, says Sachs. But they also knew it was too slow to be practical, and the metal powders required for the process were far too expensive at the time. Sachs turned to other research interests, including an effort to improve the manufacturing of photovoltaics (see "Praying for an Energy Miracle," March/April 2011). In the next decades 3-D printing took off and captured the imagination of many product designers. Most famously, a cheap and easy-to-use 3-D printer from MakerBot was introduced in 2009, appealing to many self-styled inventors and tinkerers. But these affordable printers bumped up against the reality that they were limited to using a few cheap plastics. What's more, though the machines can print complex shapes, the final product often isn't as good as a plastic part made with conventional technology.

Meanwhile, researchers at industrial manufacturers like GE were busy advancing laser-based technologies invented in the late 1980s for printing metals. These machines use lasers—or, in some cases, high-power electron beams—to draw shapes in a layer of metal powder by melt-

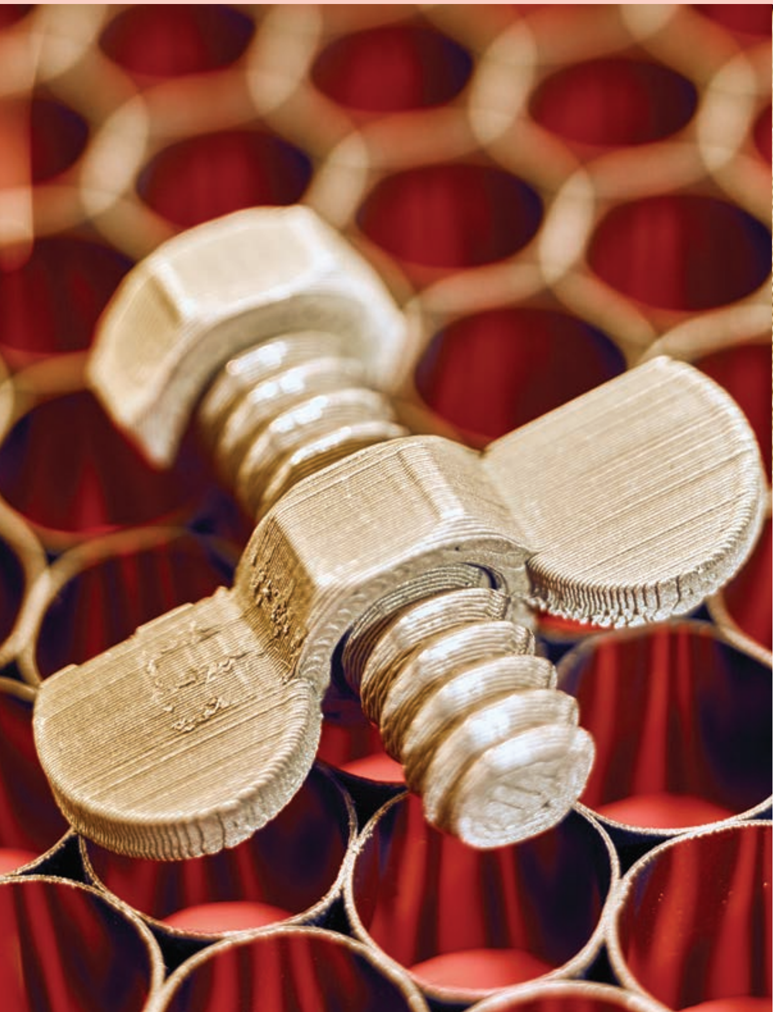

Desktop Metal printed the screw and wing nut (seen close up, above) separately to demonstrate that it can fabricate parts with tight tolerances.

ing the material. They repeat the process to build up a three-dimensional object out of the fused powders. The technique is impressive in its capabilities, but it's slow and expensive. It is worthwhile only for extremely high-value parts that are too complex to make using other methods. Notably, GE's new jet engine uses a series of sophisticated 3-D-printed fuel nozzles; they are lighter and far more durable because intricate cooling channels have been built into them.

The founders of Desktop Metal decided that to make 3-D metal printing more widely accessible, they would need to sell two different types of machines: a relatively inexpensive "desktop" model suitable for designers and engineers fabricating prototypes, and one that is fast and large enough for manufacturers. Luckily, several innovations have finally made Sachs's original invention practical for mass production, including the development of very high-speed ink-jet printing for depositing the binder. Successively printing about 1,500 layers, each 50 micrometers thick and deposited in a few seconds, the production-scale printer can build up a 500-cubic-inch part in an hour. That's about 100 times faster than a laser-based 3-D printer can make metal parts.

For its prototyping machine, Desktop Metal adopted a method from plastic-based 3-D printing. But instead of a softened polymer, it uses metal powders mixed with a flowable polymer binder. The formulation is extruded, using the printed binder to clump the metal powder into the intended shapes.

However, whether the part is printed with the prototyping machine or the production model, the resulting object—part plastic binder and part metal—lacks the strength of a metal one. So it goes into a specially designed microwave oven for sintering, a process of using heat to make the material more dense, producing a part with the desired properties. In a series of carefully calibrated steps during the sintering process, the polymer is burned off, and then the metal is fused together at a temperature well below its melting point.

The sales pitch

According to the promises of its enthusiasts, 3-D printing will reduce the need for industrial manufacturers and empower local artisan producers (see "The Difference Between Makers and Manufacturers,"

KEY PLAYERS IN 3-D PRINTING

COMPANY	TECHNOLOGY	PRODUCTS
Stratasys	One of the original 3-D-printing companies, Stratasys was founded by Scott Crumb, the inventor of fused deposition modeling, the most common way to print plastic parts.	Sells machines that can print a variety of photopolymer and thermoplastic materials.
Carbon	This Silicon Valley startup has developed a novel photochemical process for fabricating parts out of various plastics, including polyurethane and epoxy.	Introduced a modular system for manufacturers this spring.
HP	Its line of machines exploits the company's long history with ink-jet printing through what it calls "multi jet fusion technology." This uses multiple nozzles for high-speed and high-resolution printing.	Introduced its first 3-D printers last year. The initial machines print nylon, but the company is looking to expand to other materials.
3D Systems	The first 3-D-printing company, 3D Systems was founded by Chuck Hull, the inventor of stereolithography, which uses light to form parts out of photopolymers. It now offers various types of 3-D printers, including some that print metal parts.	Introduced the latest iteration of stereolithography last year.

January/February 2013). The reality is likely to be far different but nonetheless profound. Many sectors of industrial production increasingly use automation and advanced software, and 3-D printing enhances this ongoing move to digital manufacturing. In some ways, it is not unlike an automated machining process that works off a digital file to create a metal part. What's different about 3-D printing is that it offers ways to make far more complex objects and removes many of the constraints that the production process puts on designers and engineers.

It could also inspire manufacturers to change their logistics and production strategies. For relatively small quantities of goods, 3-D printing could be cheaper, since it eliminates the costs associated with the tooling, casting, and molds required to churn out most metal and plastic objects. The time and money needed to set all that up is one reason why mass production is often required if a manufacturer is going to make money. Without that incentive to commit to mass-scale production, factories could shift production schedules and be more responsive to demand, moving even closer to just-in-time manufacturing. John Hart, a professor of mechanical engineering at MIT and cofounder of Desktop Metal, calls it customized mass production. Rather than having large facilities make a huge number of identical parts that have to be shipped across the world and warehoused, manufacturers might maintain scattered factories that make a diverse set of products, ramping up production as needed. "The implications in a decade or two are probably beyond our imagination," Hart says. "I don't really think we know what we will do with these technologies."

For now, the challenge for Desktop Metal is to get its equipment in the hands of designers and engineers who are responsible for their companies' next generation of products. This winter Fulop was preparing to showcase the company's initial product, the prototyping machine, at a trade show in Pittsburgh in early May. (The production 3-D printer is scheduled to be available next year.) His task would be to convince attendees that spending $120,000 on Desktop Metal's prototyping printer and sintering furnace is essential for the future of their companies.

It is a sales job that Fulop is well suited for. He has started more than a half-dozen companies, beginning with one that imported computer hardware and software that he founded when he was 16 and still living in his native Venezuela. He is probably best known for founding A123 Systems, a battery company that was one of the highest-flying startups in the late 2000s, culminating with a $371 million IPO in 2009. The company was based on a novel lithium-ion technology developed by Yet-Ming Chiang, an MIT professor who is also a cofounder of Desktop Metal. Like their current 3-D-printing startup, A123 hoped to apply materials science expertise to revolutionize a huge market.

Though A123 enjoyed rapid growth and a highly successful IPO, the company declared bankruptcy in 2012 (Fulop left in 2010). Ask Fulop the lesson from A123 and he says simply: "Batteries are a low-margin market." Indeed, A123 struggled to compete in an increasingly crowded battery business, and it didn't offer a radical enough performance improvement over established lithium-ion batteries to immediately win over a fledging hybrid-vehicles market (see "A123's Technology Just Wasn't Good Enough" at TechnologyReview.com).

The challenges faced by Desktop Metal will be very different. A huge market for metal parts already exists. And the startup believes its technology will, at least in the short run, have few direct competitors. Chiang points to the startup's "really rich" patent portfolio. "It's not just the materials; it's the techniques, it's the [sintering] furnace," he says. "The harder the technology is, the higher the barrier to entry you build if you're successful."

In his office, Chiang has a wooden box containing a half-dozen swords, on loan from the Museum of Fine Arts in Boston, that were made in the 1970s using traditional Japanese techniques. Chiang uses the swords in teaching. The lesson: how the craftsmen used the secrets of metallurgy to turn iron ore into the final product—an ultra-sharp, slightly curved steel sword. Showing off the swords, Chiang points to some of their details, explaining the tricks their makers used, such as the quenching method used to create an extremely hard edge and a softer body. Back at his desk, his attention again on Desktop Metal, he's equally enthusiastic as he describes the metal objects recently printed by the company and on display at its facilities. What's exciting is "the idea that you can really make these parts," Chiang says. "A few hours, and here's a part that you couldn't even make before."

It won't replace such century-old production techniques as forging and metal casting, but 3-D printing could create new possibilities in manufacturing—and, just maybe, reimagine the art of metallurgy.

David Rotman is the editor of MIT Technology Review.

One of the key advantages of 3-D printing is its ability to make complex structures, including internal lattices in a metal part. Such structures could be used to make lighter and stronger parts.

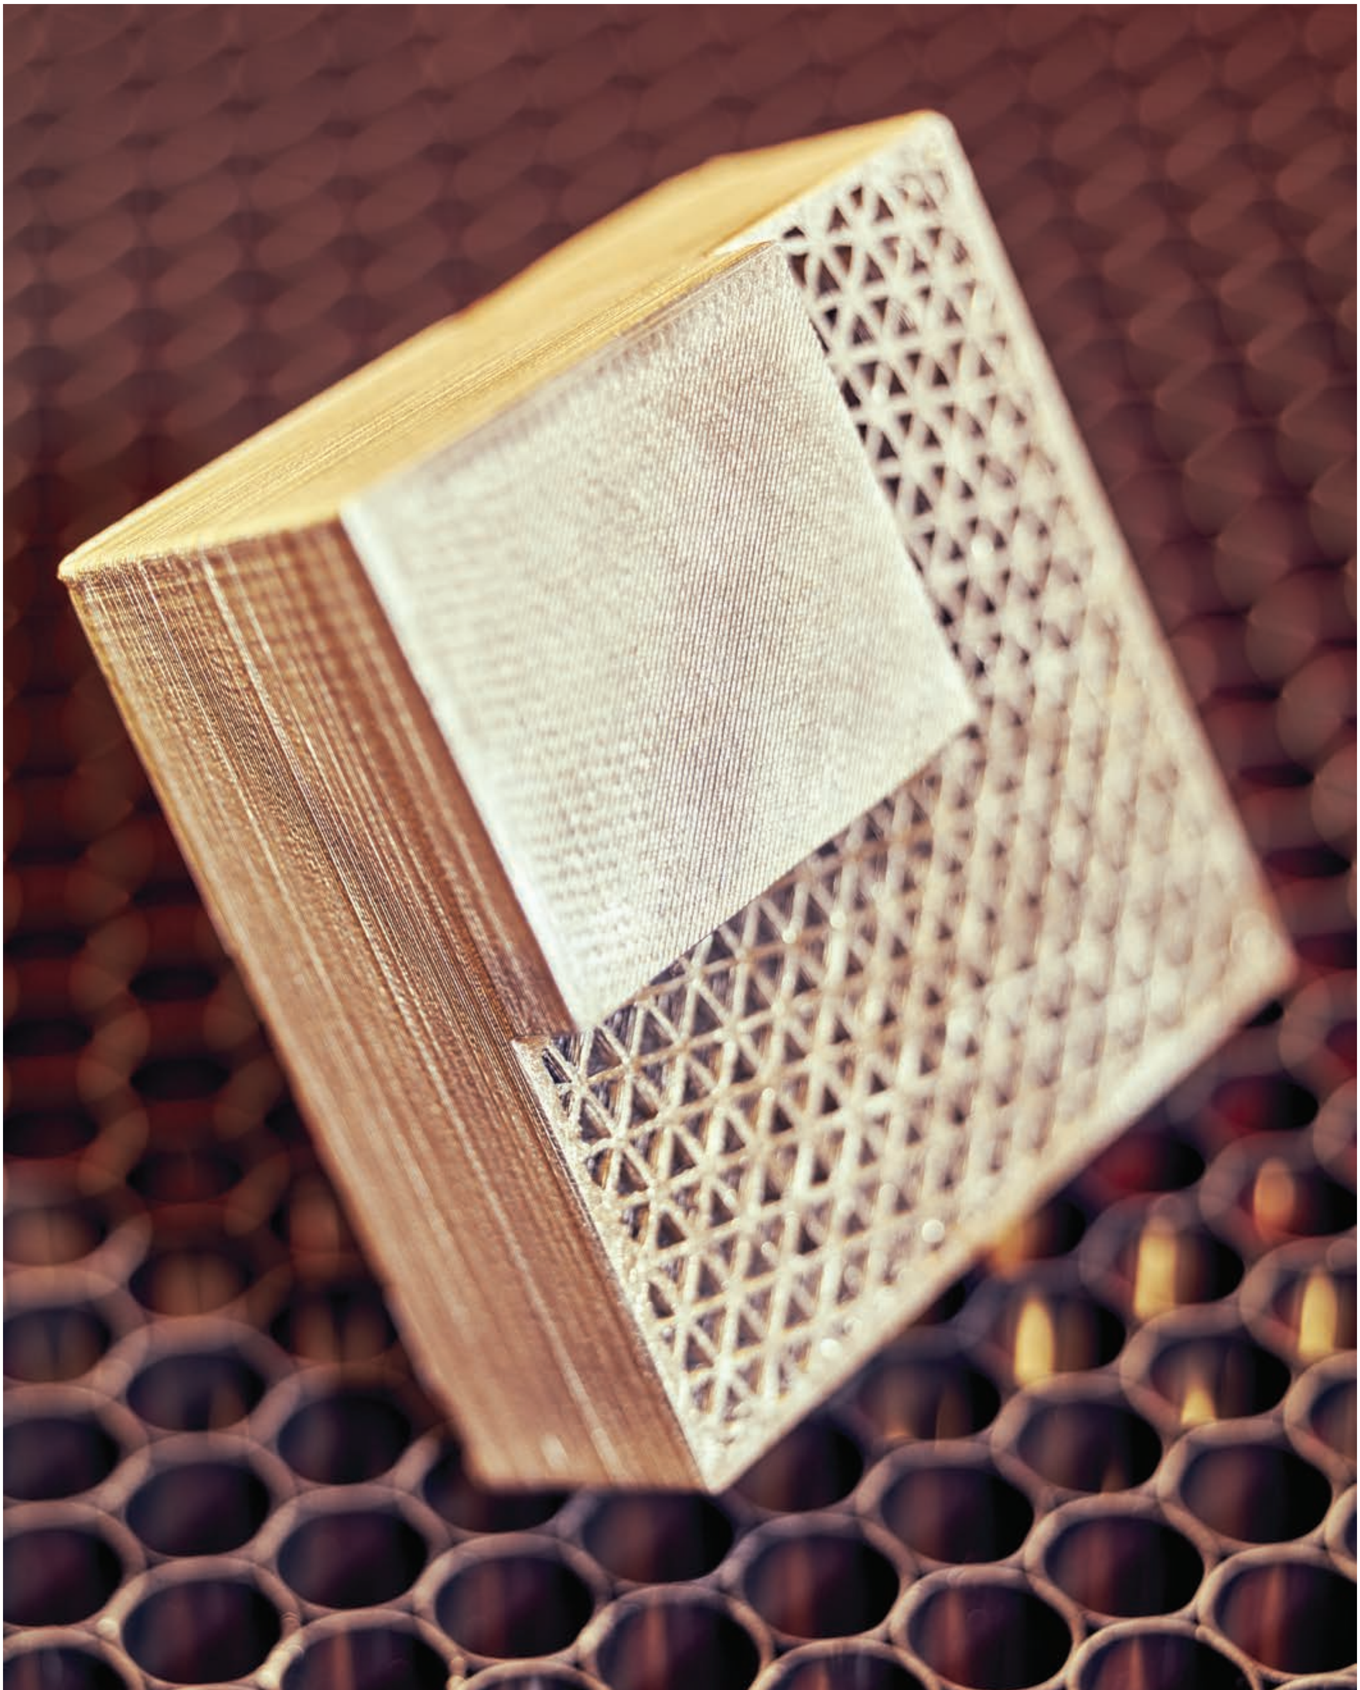

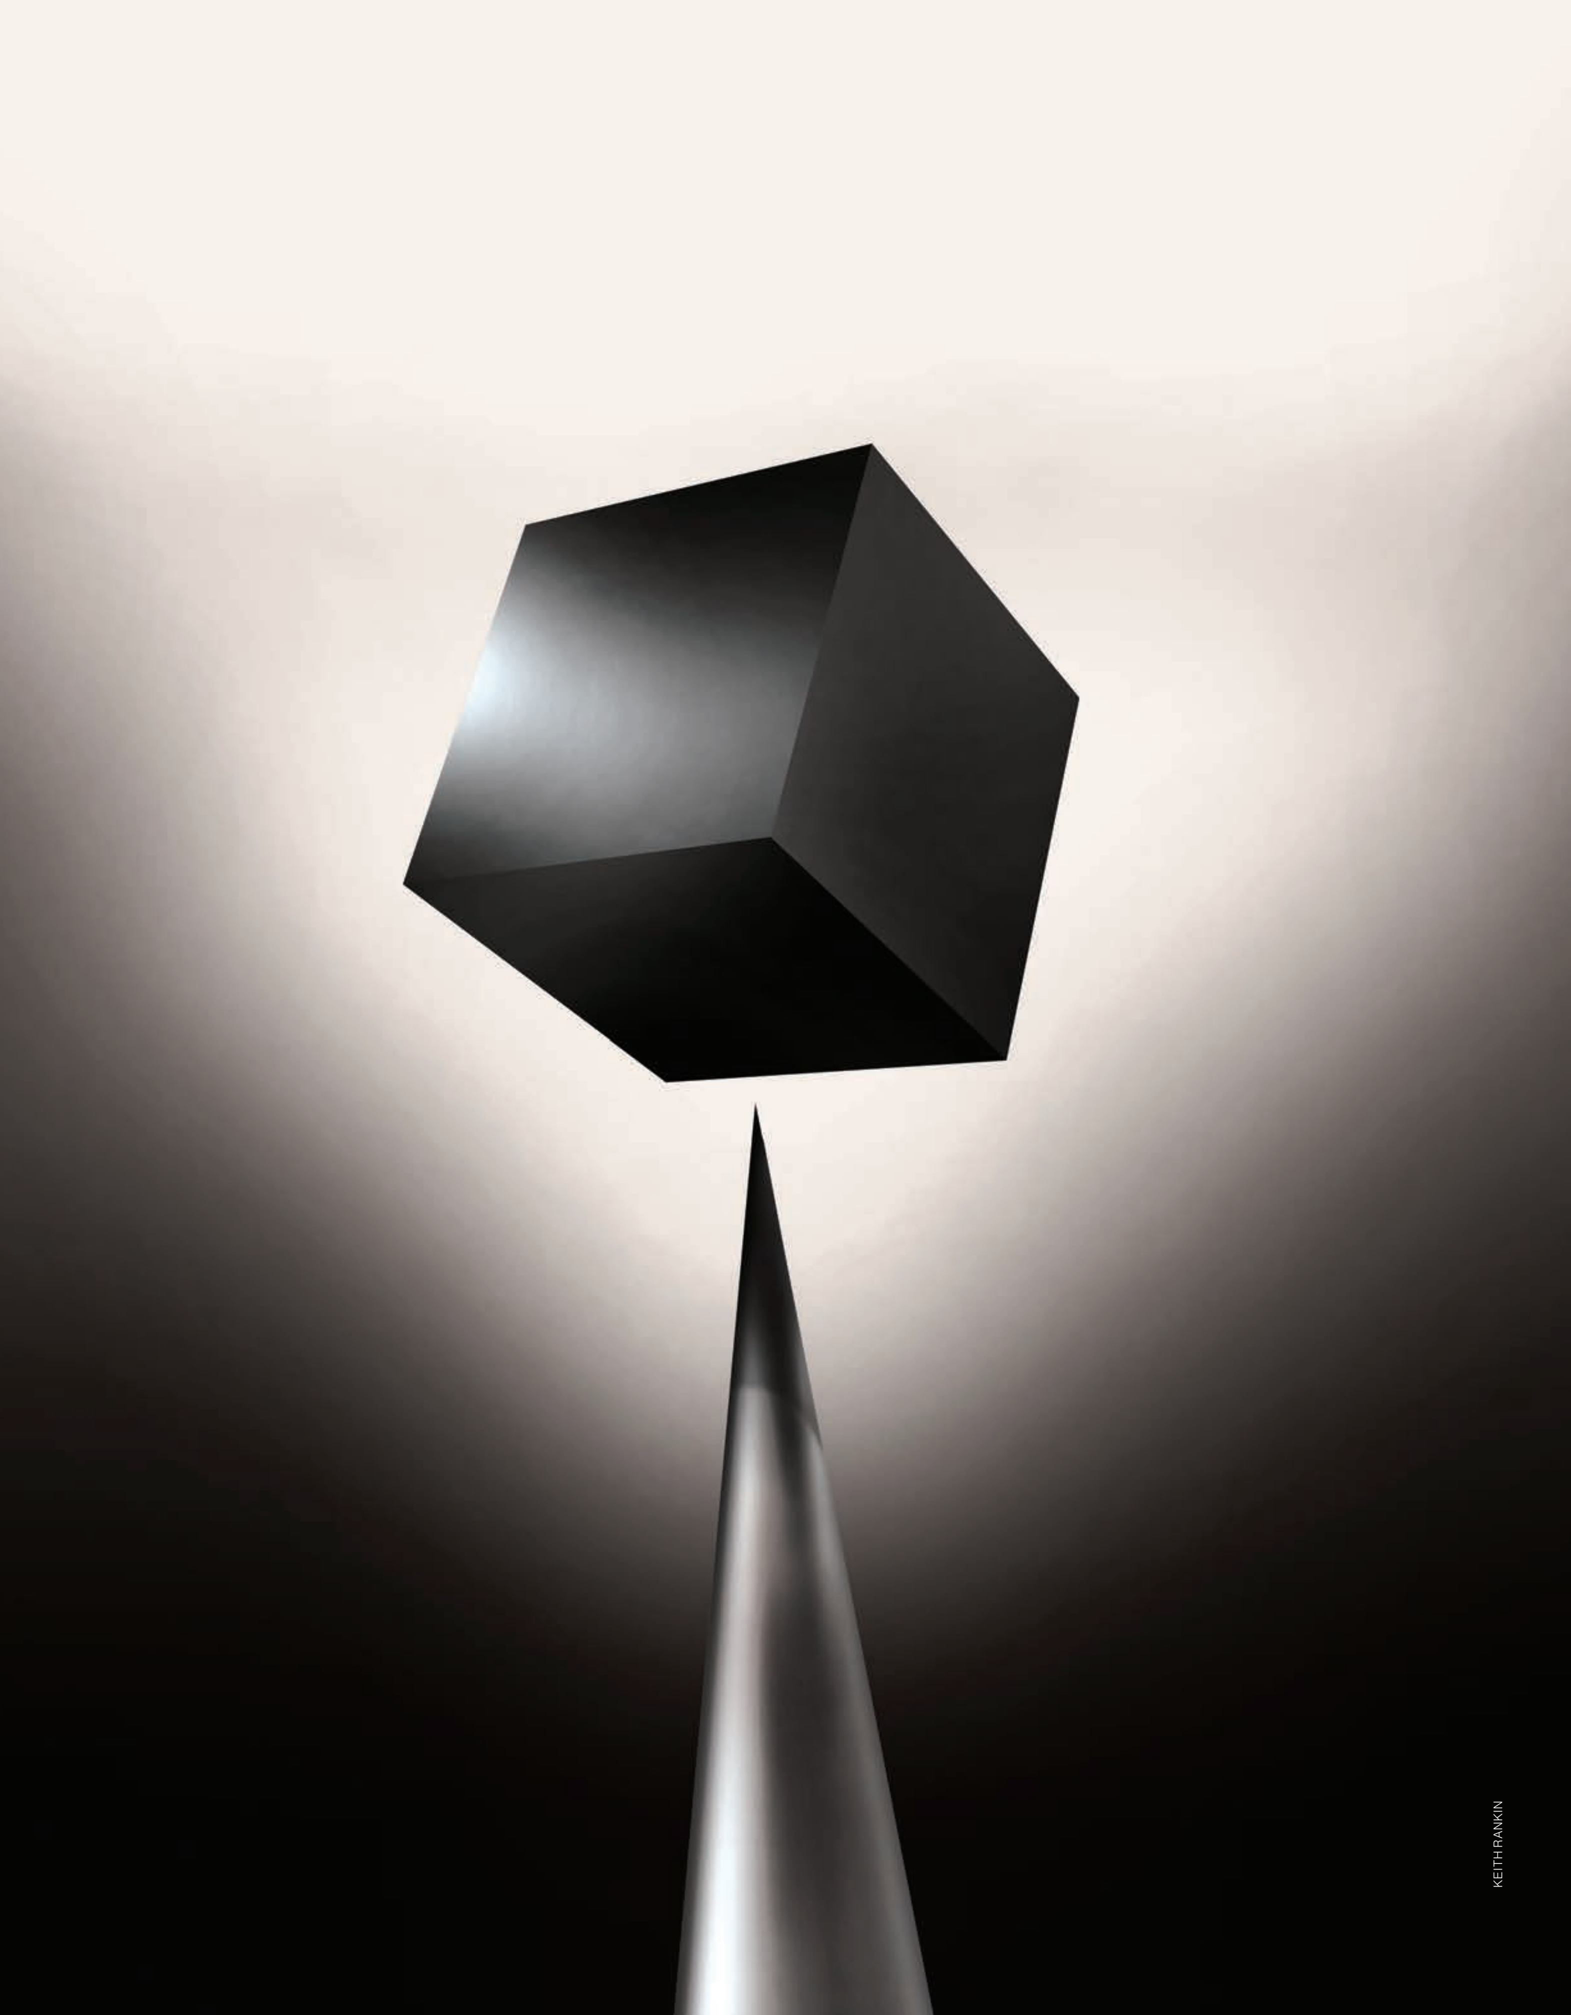

KEITH RANKIN

The Dark Secret at the Heart of AI

NO ONE REALLY KNOWS HOW THE MOST ADVANCED ALGORITHMS DO WHAT THEY DO. THAT COULD BE A SERIOUS PROBLEM AS COMPUTERS BECOME MORE RESPONSIBLE FOR MAKING IMPORTANT DECISIONS.

BY WILL KNIGHT

"WE CAN BUILD THESE MODELS,

Last year, a strange self-driving car was released onto the quiet roads of Monmouth County, New Jersey. The experimental vehicle, developed by researchers at the chip maker Nvidia, didn't look different from other autonomous cars, but it was unlike anything demonstrated by Google, Tesla, or General Motors, and it showed the rising power of artificial intelligence. The car didn't follow a single instruction provided by an engineer or programmer. Instead, it relied entirely on an algorithm that had taught itself to drive by watching a human do it.

Getting a car to drive this way was an impressive feat. But it's also a bit unsettling, since it isn't completely clear how the car makes its decisions. Information from the vehicle's sensors goes straight into a huge network of artificial neurons that process the data and then deliver the commands required to operate the steering wheel, the brakes, and other systems. The result seems to match the responses you'd expect from a human driver. But what if one day it did something unexpected—crashed into a tree, or sat at a green light? As things stand now, it might be difficult to find out why. The system is so complicated that even the engineers who designed it may struggle to isolate the reason for any single action. And you can't ask it: there is no obvious way to design such a system so that it could always explain why it did what it did.

The mysterious mind of this vehicle points to a looming issue with artificial intelligence. The car's underlying AI technology, known as deep learning, has proved very powerful at solving problems in recent years, and it has been widely deployed for tasks like image captioning, voice recognition, and language translation. There is now hope that the same techniques will be able to diagnose deadly diseases, make million-dollar trading decisions, and do countless other things to transform whole industries.

But this won't happen—or shouldn't happen—unless we find ways of making techniques like deep learning more understandable to their creators and accountable to their users. Otherwise it will be hard to predict when failures might occur—and it's inevitable they will. That's one reason Nvidia's car is still experimental.

Already, mathematical models are being used to help determine who makes parole, who's approved for a loan, and who gets hired for a job. If you could get access to these mathematical models, it would be possible to understand their reasoning. But banks, the military, employers, and others are now turning their attention to more complex machine-learning approaches that could make automated decision-making altogether inscrutable. Deep learning, the most common of these approaches, represents a fundamentally different way to program computers. "It is a problem that is already relevant, and it's going to be much more relevant in the future," says Tommi Jaakkola, a professor at MIT who works on applications of machine learning. "Whether it's an investment decision, a medical decision, or maybe a military decision, you don't want to just rely on a 'black box' method."

There's already an argument that being able to interrogate an AI system about how it reached its conclusions is a fundamental legal right. Starting in the summer of 2018, the European Union may require that companies be able to give users an explanation for decisions that automated systems reach. This might be impossible, even for systems that seem relatively simple on the surface, such as the apps and websites that use deep learning to serve ads or recommend songs. The computers that run those services have programmed themselves, and they have done it in ways we cannot understand. Even the engineers who build these apps cannot fully explain their behavior.

This raises mind-boggling questions. As the technology advances, we might soon cross some threshold beyond which using AI requires a leap of faith. Sure, we humans can't always truly explain our thought processes either—but we find ways to intuitively trust and gauge people. Will that also be possible with machines that think and make decisions differently from the way a human would? We've never before built machines that operate in ways their creators don't understand. How well can we expect to communicate—and get along with—intelligent machines that could be unpredictable and inscrutable? These questions took me on a journey to the bleeding edge of research on AI algorithms, from Google to Apple and many places in between, including a meeting with one of the great philosophers of our time.

Doctor bot

In 2015, a research group at Mount Sinai Hospital in New York was inspired to

apply deep learning to the hospital's vast database of patient records. This data set features hundreds of variables on patients, drawn from their test results, doctor visits, and so on. The resulting program, which the researchers named Deep Patient, was trained using data from about 700,000 individuals, and when tested on new records, it proved incredibly good at predicting disease. Without any expert instruction, Deep Patient had discovered patterns hidden in the hospital data that seemed to indicate when people were on the way to a wide range of ailments, including cancer of the liver. There are a lot of methods that are "pretty good" at predicting disease from a patient's records, says Joel Dudley, who leads the Mount Sinai team. But, he adds, "this was just way better."

At the same time, Deep Patient is a bit puzzling. It appears to anticipate the onset of psychiatric disorders like schizophrenia surprisingly well. But since schizophrenia is notoriously difficult for physicians to predict, Dudley wondered how this was possible. He still doesn't know. The new tool offers no clue as to how it does this. If something like Deep Patient is actually going to help doctors, it will ideally give them the rationale for its prediction, to reassure them that it is accurate and to justify, say, a change in the drugs someone is being prescribed. "We can build these models," Dudley says ruefully, "but we don't know how they work."

Artificial intelligence hasn't always been this way. From the outset, there were two schools of thought regarding how understandable, or explainable, AI ought to be. Many thought it made the most sense to build machines that reasoned according to rules and logic, making their inner workings transparent to anyone who cared to examine some code. Others felt that intelligence would more easily emerge if machines took inspiration from biology, and learned by observing and experiencing. This meant turning computer programming on its head. Instead of a programmer writing the commands to solve a problem, the program generates its own algorithm based on example data and a desired output. The machine-learning techniques that would later evolve into today's most powerful AI systems followed the latter path: the machine essentially programs itself.

At first this approach was of limited practical use, and in the 1960s and '70s it remained largely confined to the fringes of the field. Then the computerization of many industries and the emergence of large data sets renewed interest. That inspired the development of more powerful machine-learning techniques, especially new versions of one known as the artificial neural network. By the 1990s, neural networks could automatically digitize handwritten characters.

But it was not until the start of this decade, after several clever tweaks and refinements, that very large—or "deep"—neural networks demonstrated dramatic improvements in automated perception. Deep learning is responsible for today's explosion of AI. It has given computers extraordinary powers, like the ability to recognize spoken words almost as well as a person could, a skill too complex to code into the machine by hand. Deep learning has transformed computer vision and dramatically improved machine translation. It is now being used to guide all sorts of key decisions in medicine, finance, manufacturing—and beyond.

BUT WE DON'T KNOW HOW THEY WORK."

The workings of any machine-learning technology are inherently more opaque, even to computer scientists, than a hand-coded system. This is not to say that all future AI techniques will be equally unknowable. But by its nature, deep learning is a particularly dark black box.

You can't just look inside a deep neural network to see how it works. A network's reasoning is embedded in the behavior of thousands of simulated neurons, arranged into dozens or even hundreds of intricately interconnected layers. The neurons in the first layer each receive an input, like the intensity of a pixel in an image, and then perform a calculation before outputting a new signal. These outputs are fed, in a complex web, to the neurons in the next layer, and so on, until an overall output is produced. Plus, there is a process known as back-propagation that tweaks the calculations of individual neurons in a way that lets the network learn to produce a desired output.

The many layers in a deep network enable it to recognize things at different levels of abstraction. In a system designed


ADAM FERRISS

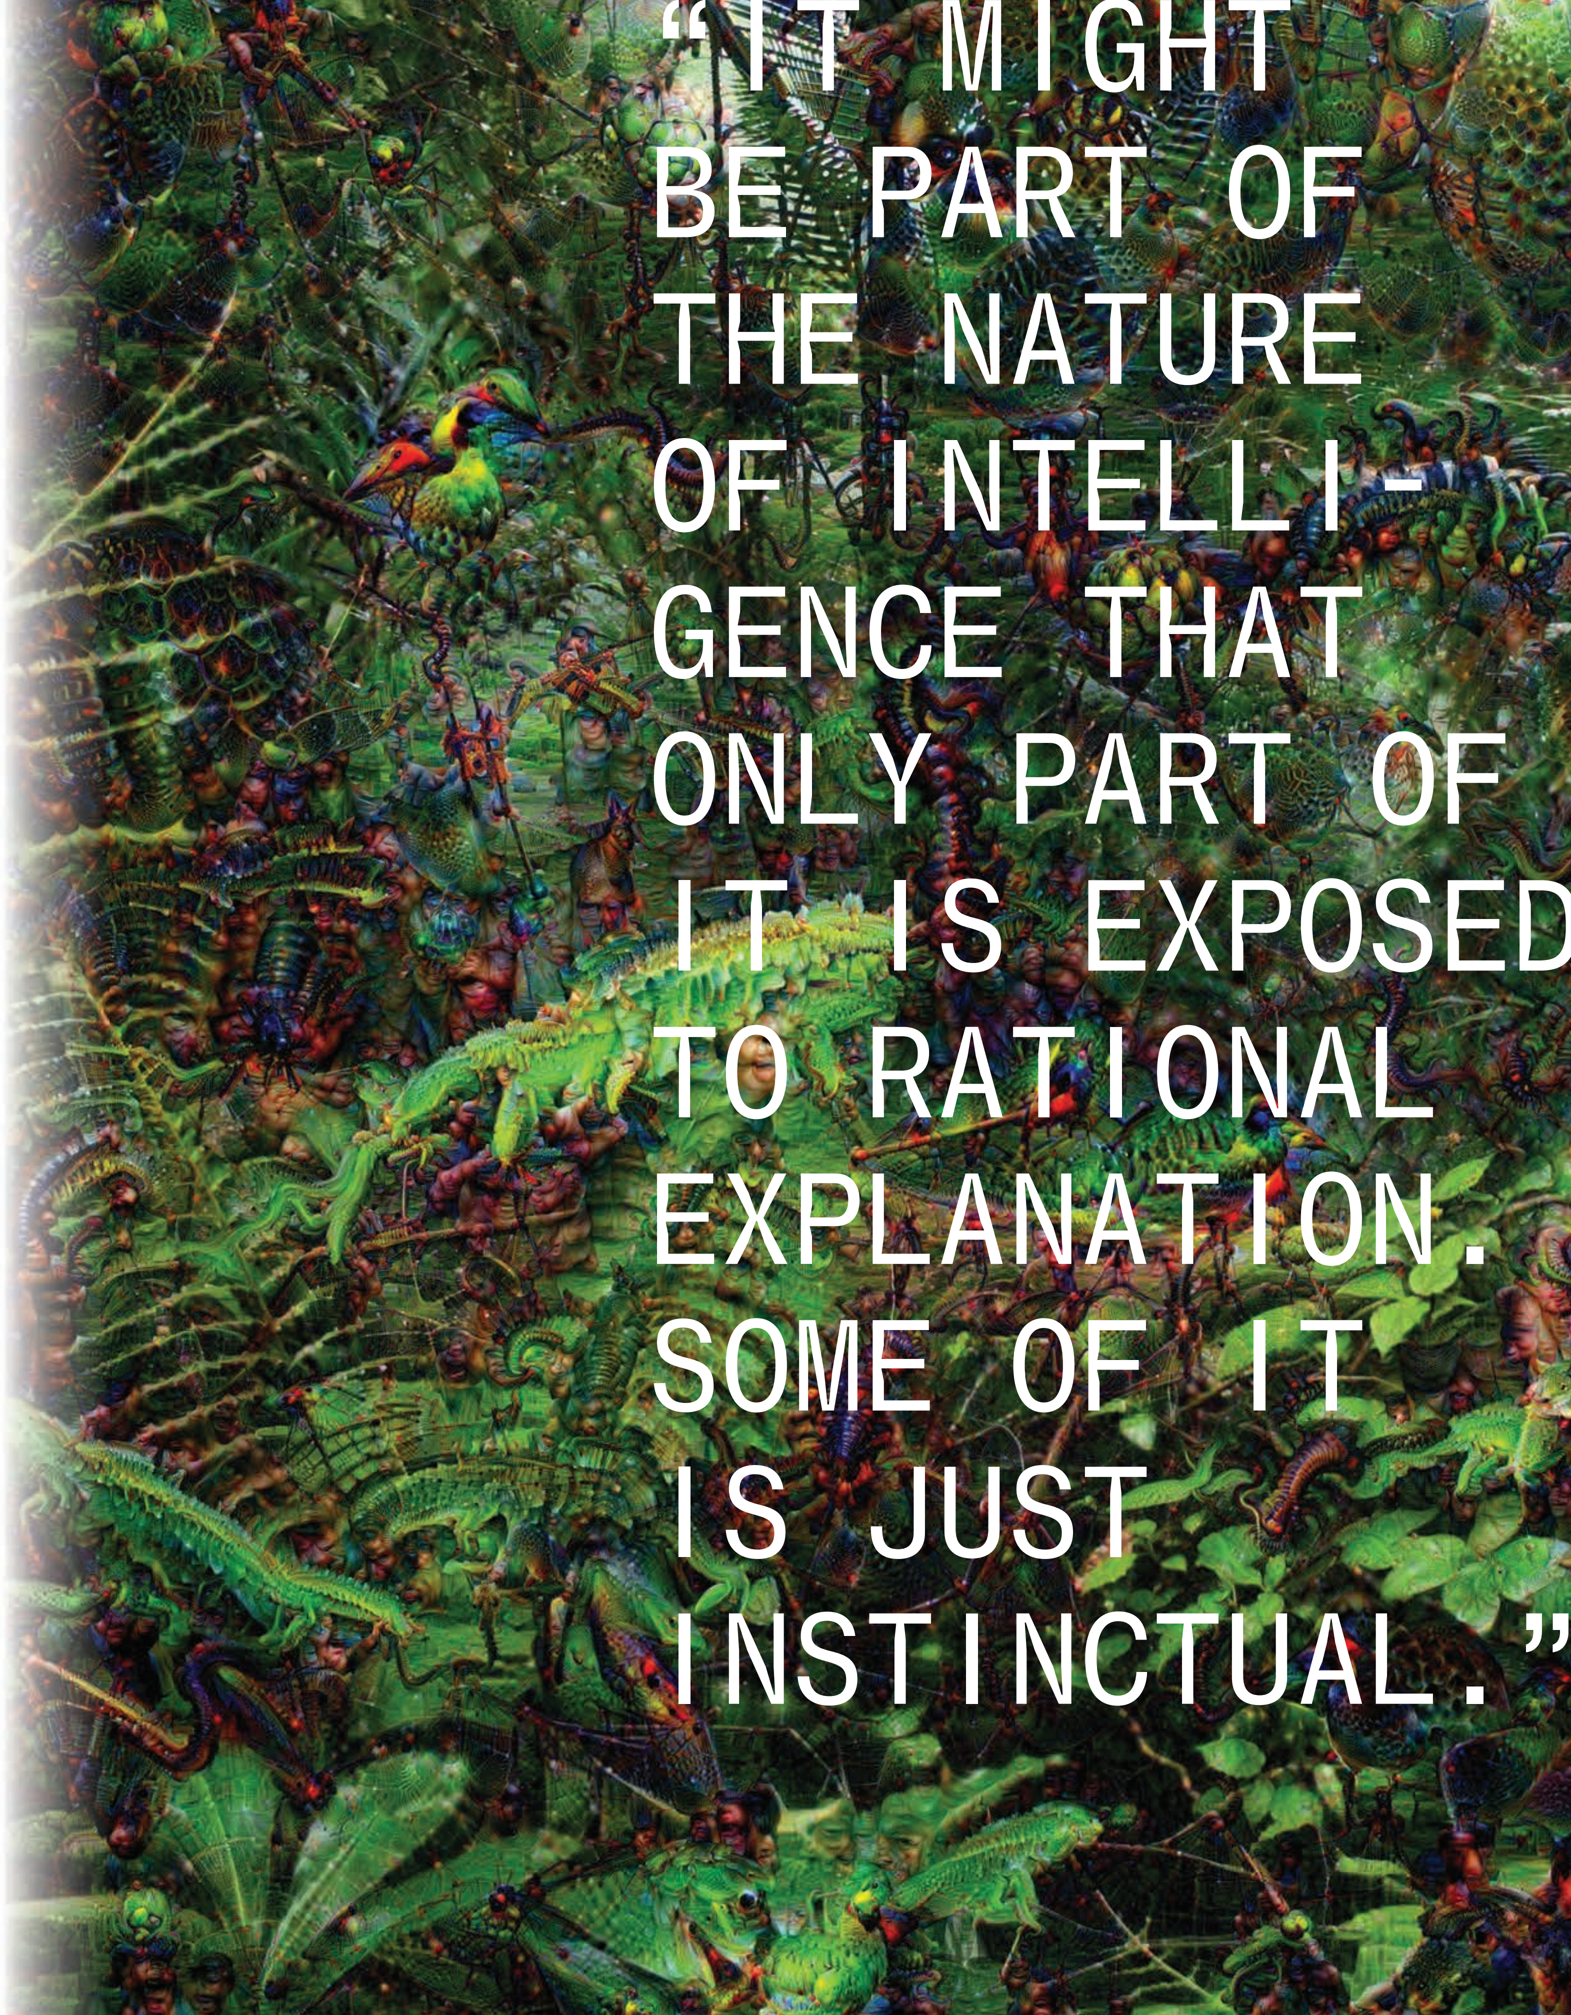

"IT MIGHT BE PART OF THE NATURE OF INTELLIGENCE THAT ONLY PART OF IT IS EXPOSED TO RATIONAL EXPLANATION. SOME OF IT IS JUST INSTINCTUAL."

HOW WELL CAN WE GET ALONG WITH MACHINES THAT ARE UNPREDICT-ABLE AND INSCRUTA-BLE?

to recognize dogs, for instance, the lower layers recognize simple things like outlines or color; higher layers recognize more complex stuff like fur or eyes; and the topmost layer identifies it all as a dog. The same approach can be applied, roughly speaking, to other inputs that lead a machine to teach itself: the sounds that make up words in speech, the letters and words that create sentences in text, or the steering-wheel movements required for driving.

Ingenious strategies have been used to try to capture and thus explain in more detail what's happening in such systems. In 2015, researchers at Google modified a deep-learning-based image recognition algorithm so that instead of spotting objects in photos, it would generate or modify them. By effectively running the algorithm in reverse, they could discover the features the program uses to recognize, say, a bird or building. The resulting images, produced by a project known as Deep Dream, showed grotesque, alien-like animals emerging from clouds and plants, and hallucinatory pagodas blooming across forests and mountain ranges. The images proved that deep learning need not be entirely inscrutable; they revealed that the algorithms home in on familiar visual features like a bird's beak or feathers. But the images also hinted at how different deep learning is from human perception, in that it might make something out of an artifact that we would know to ignore. Google researchers noted that when its algorithm generated images of a dumbbell, it also generated a human arm holding it. The machine had concluded that an arm was part of the thing.

Further progress has been made using ideas borrowed from neuroscience and cognitive science. A team led by Jeff Clune, an assistant professor at the University of Wyoming, has employed the AI equivalent of optical illusions to test deep neural networks. In 2015, Clune's group showed how certain images could fool such a network into perceiving things that aren't there, because the images exploit the low-level patterns the system searches for. One of Clune's collaborators, Jason Yosinski, also built a tool that acts like a probe stuck into a brain. His tool targets any neuron in the middle of the network and searches for the image that activates it the most. The images that turn up are abstract (imagine an impressionistic take on a flamingo or a school bus), highlighting the mysterious nature of the machine's perceptual abilities.

We need more than a glimpse of AI's thinking, however, and there is no easy solution. It is the interplay of calculations inside a deep neural network that is crucial to higher-level pattern recognition and complex decision-making, but those calculations are a quagmire of mathematical functions and variables. "If you had a very small neural network, you might be able to understand it," Jaakkola says. "But once it becomes very large, and it has thousands of units per layer and maybe hundreds of layers, then it becomes quite un-understandable."

In the office next to Jaakkola is Regina Barzilay, an MIT professor who is determined to apply machine learning to medicine. She was diagnosed with breast cancer a couple of years ago, at age 43. The diagnosis

was shocking in itself, but Barzilay was also dismayed that cutting-edge statistical and machine-learning methods were not being used to help with oncological research or to guide patient treatment. She says AI has huge potential to revolutionize medicine, but realizing that potential will mean going beyond just medical records. She envisions using more of the raw data that she says is currently underutilized: "imaging data, pathology data, all this information."

After she finished cancer treatment last year, Barzilay and her students began working with doctors at Massachusetts General Hospital to develop a system capable of mining pathology reports to identify patients with specific clinical characteristics that researchers might want to study. However, Barzilay understood that the system would need to explain its reasoning. So, together with Jaakkola and a student, she added a step: the system extracts and highlights snippets of text that are representative of a pattern it has discovered. Barzilay and her students are also developing a deep-learning algorithm capable of finding early signs of breast cancer in mammogram images, and they aim to give this system some ability to explain its reasoning, too. "You really need to have a loop where the machine and the human collaborate," Barzilay says.

Killing machines

The U.S. military is pouring billions into projects that will use machine learning to pilot vehicles and aircraft, identify targets, and help analysts sift through huge piles of intelligence data. Here more than anywhere else, even more than in medicine, there is little room for algorithmic mystery, and the Department of Defense has identified explainability as a key stumbling block.

David Gunning, a program manager at the Defense Advanced Research Projects Agency, is overseeing the aptly named Explainable Artificial Intelligence program. A silver-haired veteran of the agency who previously oversaw the DARPA project that eventually led to the creation of Siri, Gunning says automation is creeping into countless areas of the military. Intelligence analysts are testing machine learning as a way of identifying patterns in vast amounts of surveillance data. Many autonomous ground vehicles and aircraft are being developed and tested. But soldiers probably won't feel comfortable in a robotic tank that doesn't explain itself to them, and analysts will be reluctant to act on information without some reasoning. "It's often the nature of these machine-learning systems that they produce a lot of false alarms, so an intel analyst really needs extra help to understand why a recommendation was made," Gunning says.

This March, DARPA chose 13 projects from academia and industry for funding under Gunning's program. Some of them could build on work led by Carlos Guestrin, a professor at the University of Washington. He and his colleagues have developed a way for machine-learning systems to provide a rationale for their outputs. Essentially, under this method a computer automatically finds a few examples from a data set and serves them up in a short explanation. A system designed to classify an e-mail message as coming from a terrorist, for example, might use many millions of messages in its training and decision-making. But using the Washington team's approach, it could highlight certain keywords found in a message. Guestrin's group has also devised ways for image recognition systems to hint at their reasoning by highlighting the parts of an image that were most significant.

One drawback to this approach and others like it, such as Barzilay's, is that the explanations provided will always be simplified, meaning some vital information may be lost along the way. "We haven't achieved the whole dream, which is where AI has a conversation with you, and it is able to explain," says Guestrin. "We're a long way from having truly interpretable AI."

It doesn't have to be a high-stakes situation like cancer diagnosis or military maneuvers for this to become an issue. Knowing AI's reasoning is also going to be crucial if the technology is to become a common and useful part of our daily lives. Tom Gruber, who leads the Siri team at Apple, says explainability is a key consideration for his team as it tries to make Siri a smarter and more capable virtual assistant. Gruber wouldn't discuss specific plans for Siri's future, but it's easy to imagine that if you receive a restaurant recommendation from Siri, you'll want to know what the reasoning was. Ruslan Salakhutdinov, director of AI research at Apple and an associate professor at Carnegie Mellon University, sees explainability as the core of the evolving relationship between humans and intelligent machines. "It's going to introduce trust," he says.

Going deep

Just as many aspects of human behavior are impossible to explain in detail, perhaps it won't be possible for AI to explain everything it does. "Even if somebody can give you a reasonable-sounding explanation [for his or her actions], it probably is incomplete, and the same could very well be true for AI," says Clune, of the University of Wyoming. "It might just be part of the nature of intelligence that only part of it is exposed to rational explanation. Some of it is just instinctual, or subconscious, or inscrutable."

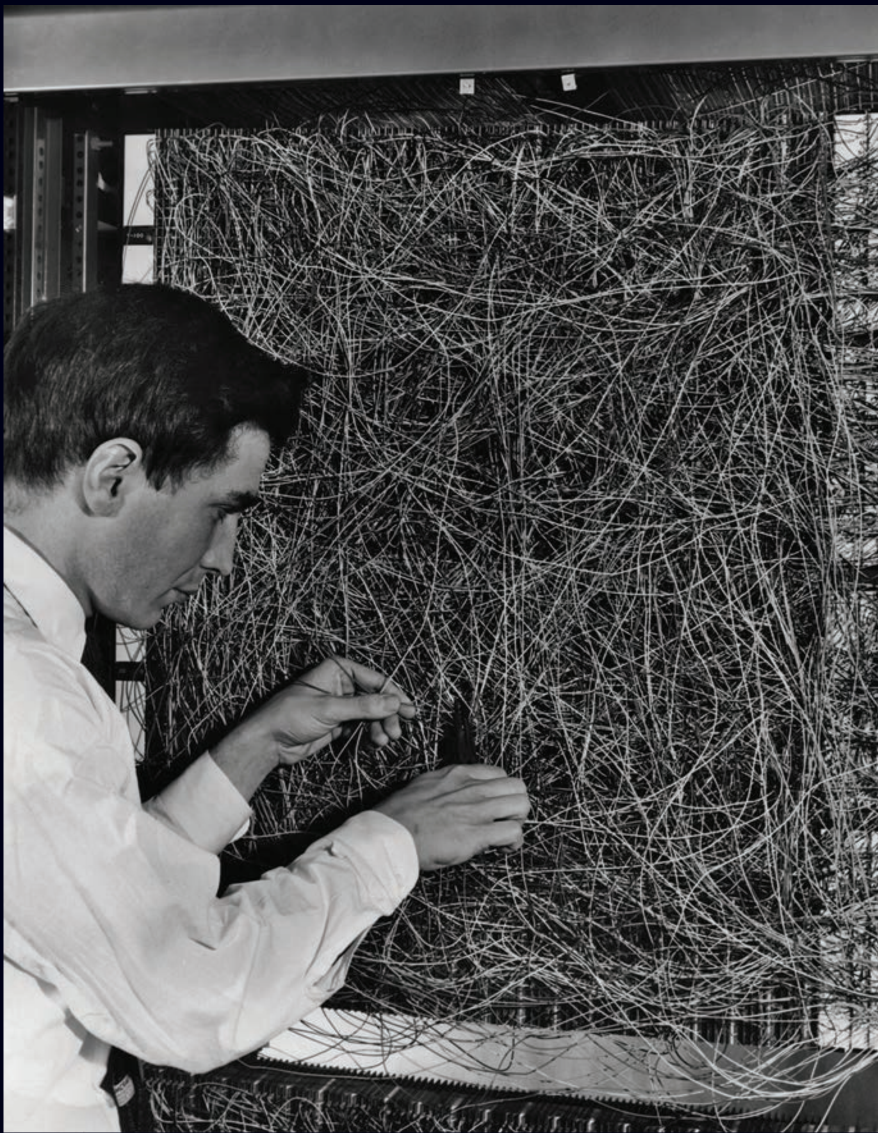

This early artificial neural network, at the Cornell Aeronautical Laboratory in Buffalo, New York, circa 1960, processed inputs from light sensors ...

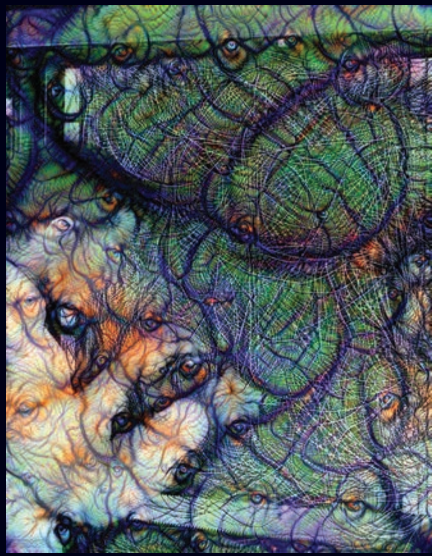

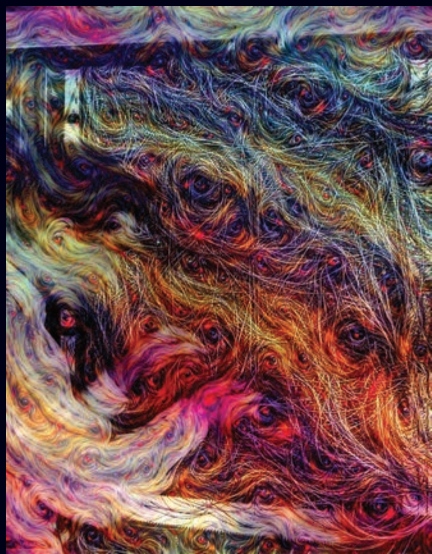

... inspiring the artist Adam Ferriss to run that image through Google Deep Dream, a progam that adjusts an image to stimulate the pattern recognition capabilities of a deep neural network. These pictures, and those on pages 58 and 59, were produced using a mid-level layer of the neural network.

ADAM FERRISS; FREDERIC LEWIS/GETTY

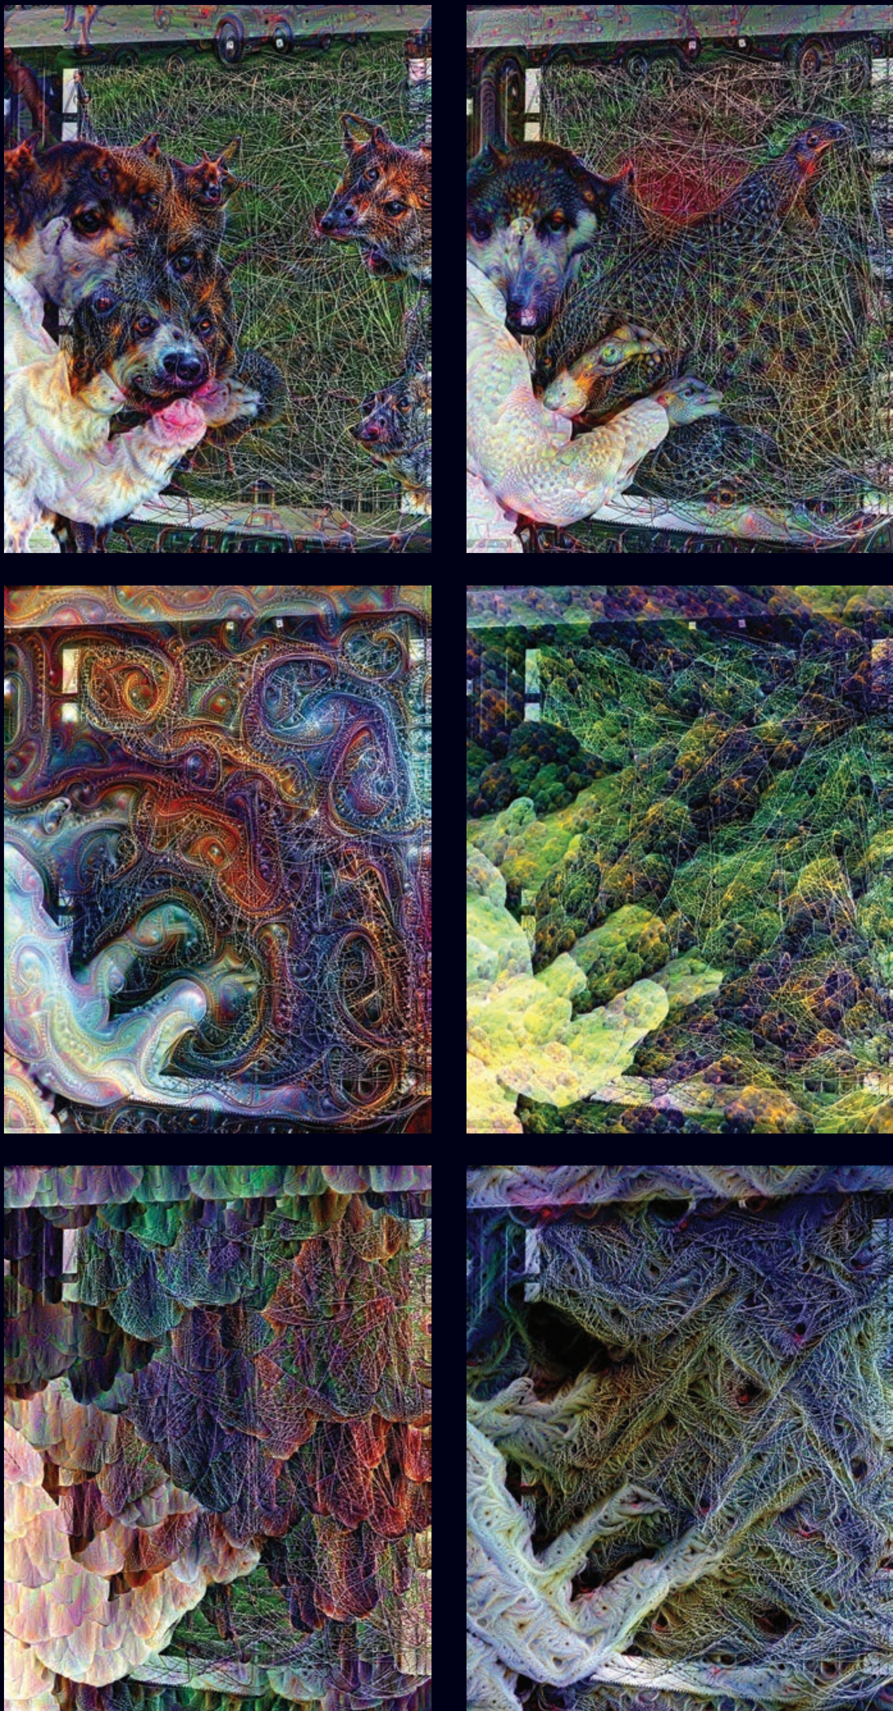

If that's so, then at some stage we may have to simply trust AI's judgment or do without using it. Likewise, that judgment will have to incorporate social intelligence. Just as society is built upon a contract of expected behavior, we will need to design AI systems to respect and fit with our social norms. If we are to create robot tanks and other killing machines, it is important that their decision-making be consistent with our ethical judgments.

To probe these metaphysical concepts, I went to Tufts University to meet with Daniel Dennett, a renowned philosopher and cognitive scientist who studies consciousness and the mind. A chapter of Dennett's latest book, *From Bacteria to Bach and Back*, an encyclopedic treatise on consciousness, suggests that a natural part of the evolution of intelligence itself is the creation of systems capable of performing tasks their creators do not know how to do. "The question is, what accommodations do we have to make to do this wisely—what standards do we demand of them, and of ourselves?" he tells me in his cluttered office on the university's idyllic campus.

He also has a word of warning about the quest for explainability. "I think by all means if we're going to use these things and rely on them, then let's get as firm a grip on how and why they're giving us the answers as possible," he says. But since there may be no perfect answer, we should be as cautious of AI explanations as we are of each other's—no matter how clever a machine seems. "If it can't do better than us at explaining what it's doing," he says, "then don't trust it."

Will Knight is MIT Technology Review*'s senior editor for artificial intelligence and robotics.*

sinforma

Disinformation

Disinformation

Disinformation

By John Pollock

Russia's reinvention of war exploits old techniques for a new century. Open-source citizen investigators are fighting back.

The biggest crime scene in the world

Members of the Ukrainian State Emergency Service search for bodies in a field at the crash site of Malaysia Airlines Flight MH17, near the village of Hrabove (Grabove) in the Donetsk region, on July 26, 2014.

BULENT KILIC/GETTY

ON JULY 17, 2014, AS PASSENGERS CHECKED IN AT AMSTERDAM'S SCHIPHOL AIRPORT FOR MALAYSIA AIRLINES FLIGHT MH17, "NECRO MANCER" (@666_mancer) tweeted about an unusual convoy 1,500 miles east in Ukraine. His citizen intelligence network had noticed a covered antiaircraft missile system trundling through Donetsk on a low loader. Minutes later, half a world away in Brasilia, Vladimir Putin wrapped up a pre-dawn Russian press conference. His answer to the last question—about the Moscow Metro's worst accident, two days earlier, which killed 24 people—was overlooked until Russian conspiracy theorists picked it up 15 months later. In the aftermath of the deadliest shoot-down in history, his words acquired a macabre resonance.

"Responsibility should always be personal," said Putin. "There is a classic example from criminal law called a 'shooting tragedy,' when two hunters shoot at a bush thinking there is game there, and accidentally kill a man. Since experts cannot establish who did it, they are both set free ... Investigators should expose the guilty party ... and they should be held responsible, but only those specific individuals whose fault it was." By day's end, questions about guilt and responsibility for MH17's downing were of global concern. International investigators would range from the official—well resourced and highly trained—to a self-organized group called Bellingcat. Using little more than laptops, open-source materials, and relentless dedication, these "citizen investigative journalists" would find the exact missile launcher, identify dozens of soldiers, and, eventually, reveal a senior Russian soldier with a key role in coördinating the missile launch.

The horror of the MH17 atrocity was peculiarly intimate: in the debris fields, stuffed toys spilled from children's suit-

ROB STOTHARD/GETTY

The Bloodlands are wetted again

Smoke rises from a building damaged during fighting between Russian-backed separatists and Ukrainian government troops on July 21, 2014, in Donetsk, Ukraine.

cases. It briefly brought Ukraine's war into focus in a way that Russia's misdirection over the annexation of Crimea, or their murky fight in the farthest corner of Europe, had failed to do. However, a deeper and wider war remained concealed.

Andrei Illarionov was Putin's senior economic advisor, and personal representative to the G8 for five years, until he resigned in protest at growing corruption. Two months before the downing of MH17, he called Ukraine "an introductory chapter" in "the Fourth World War." (Stalin's dismal term for the Cold War was "World War Three.") Illarionov dislikes the phrase, but he says it's being used "by the Kremlin propaganda machine" for a war "being waged now by Russia against the rest of the world."

It took the weaponization of information in the 2016 U.S. presidential election for the Western world to start to notice. We now know of e-mails stolen from the Democratic National Committee by Russian hackers, of sophisticated botnets, of similar attacks across Europe; but the full extent of Russia's activities is still being uncovered. Realizing that we are at war, and understanding how we can fight back, is now urgent business. The story of MH17, and Russia's exposure, offers a grim but useful case study.

DEVICES OF DISINFORMATION

As with the Soviets' shooting down of Korean Air Lines Flight 007 in 1983, when 269 died, Russia's reaction to the international outrage over MH17 was to contrive and deliver counternarratives. A buzzing and growing cloud of ever-changing claims emerged—placing the blame on a Ukrainian fighter jet, Ukrainians on the ground, or the CIA, or claiming that Putin's private plane was the real target. Russia's tactics, says Ben Nimmo,

senior fellow in information defense at the Atlantic Council's Digital Forensic Research Lab, rely on what he calls "the 4D approach": "Dismiss, distort, distract, dismay. Never confess, never admit—just keep on attacking."

"The single most prevalent Russian response is to attack the critic," he says. "They use a 'vilify and amplify' technique." Critics are besmirched, sometimes in an official announcement, sometimes through proxies, sometimes through anonymous sources quoted in state media; then paid trolls and highly automated networks of bots add scale. In response, an ad hoc blend of civilians, private companies, and NGOs has evolved to cast a bright, shining light on MH17 and Russian aggression in Ukraine, Syria, and the Atlantic partnership. Exemplifying the values Italo Calvino outlined in *Six Memos for the Next Millennium*—lightness, quickness, exactitude, visibility, multiplicity, and consistency—their methods are in sharp contrast to the West's generally sclerotic response to a revanchist Russia.

Nowhere is this weakness more brutally apparent than in Russia's use of digital technology to reinforce its greatest tool of statecraft: *maskirovka*. The literal translation—"little masquerade"—disguises the density and importance of this elusive concept. "Military deception" misses its deep cultural roots: *maskirovka* involves camouflage, denial, and a deep finesse. As James Jesus Angleton, the founding counterintelligence chief of the CIA, put it, "The myriad stratagems, deceptions, artifices, and all the other devices of disinformation ... confuse and split the West [with] an ever-fluid landscape, where fact and illusion merge, a kind of wilderness of mirrors."

The most powerful weapon in the *maskirovka* armory is disinformation, a word acquired in the 1950s from the Russian *dezinformatsiya*. A generation after the Cold War, the acknowledged masters of "deza" are deploying disinformation technology against the compromised immune system of liberal democracy. "And at this point," says Andrew Andersen, a Russian-born security analyst at the University of Calgary's Centre for Military and Strategic Studies, "the West is losing."

"The first thing you need to understand is that this *is* a war," says Andersen. "This is not a joke and not a game of any kind. It's not 'socializing with your friends on social networks'—it's a real war. Even those who don't want to take part have to behave in accordance with the laws of war," he says, alluding to Trotsky's notorious epigram, recalled by several of the interviewees for this story, that translates loosely as: "You may not be interested in war, but war is interested in you."

Even the name of this new style of war is contested territory. "Ambiguous," "hybrid," "irregular," and "nonlinear" warfare have all been suggested. Mark Galeotti, senior research fellow at Prague's Institute of International Relations, unpicks the issue in his new book, *Hybrid War or Gibridnaya Voina? Getting Russia's Non-Linear Military Challenge Right*. He admits to still worrying at it. "The more I think about what we should be calling hybrid war," he says, "the more I think the answer is: war."

"The Russians have stumbled on how the nature of international contestation is changing and will be fought out in the 21st century. It's an age when direct kinetic warfare [the military's term of art for 'things that go bang'] is ridiculously expensive, in political but also economic terms," he says. "Instead, war will be fought out through a variety of other means, many which are covert, ambiguous, and so on. The Russians have, fortuitously for them, simply stumbled on a truth of the century."

> Disinformation confuses and splits the West with "an ever-fluid landscape ... a kind of wilderness of mirrors."

FOUL DEEDS WILL RISE

Exactly an hour before MH17 took off, Necro Mancer tweeted a tentative identification: "It visually resembles a BUK

Irregular forces

During the 2014 Crimean crisis, soldiers at Ukraine's Perevalne military base displayed no insignia or badges of rank, but they included Russian troops.

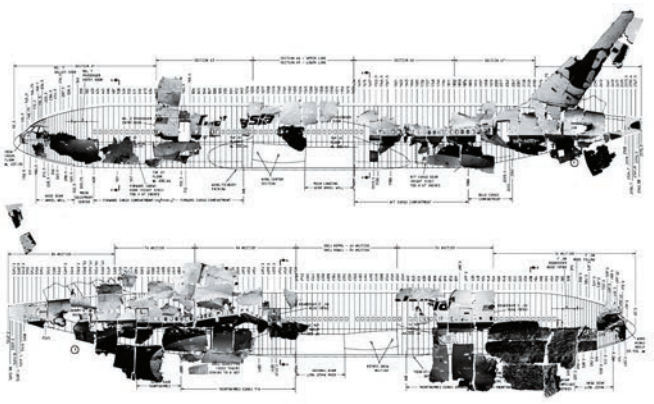

Recovered wreckage

The Dutch Safety Board was able to reconstruct a significant amount of the downed Malaysian Airlines jet.

ANTON HOLOBORODKO; DUTCH SAFETY BOARD

a lot." (Buks are a family of Russian-made mobile medium-range surface-to-air missile systems.) A Donetsk man of around 50, he spends "almost all" his free time scanning popular Russian-language social-media sites like Vkontakte ("In Contact"), known as "Russia's Facebook," and Odnoklassniki ("Classmates"); listening to pro-Russian channels on the walkie-talkie app Zello; and sharing civilian reports of military activities. As an additional hobby, he uses open sources to curate a list, linked in his Twitter profile, of several thousand Russian and pro-Russian dead, "because they hide them."

"I cannot fight as a soldier, so I try to do my best," he says of a conflict that has led to more than 30,000 casualties and millions displaced. He's just one of many

Dying bodies fell “like confetti” for 90 seconds

A photo taken July 17, 2014, shows the wreckage of MH17 near the town of Shakhtarsk, Ukraine. The plane was carrying 298 people from Amsterdam to Kuala Lumpur. Russian-backed separatists and Russian sources made various misleading claims, including that the Malaysian airliner had been shot down by a Ukrainian jet.

DOMINIQUE FAGET/GETTY

"You may not be interested in war,
but war is interested in you."

Fighting back

Eliot Higgins, a British investigative journalist, founded Bellingcat so that citizen journalists could investigate current events like the shoot-down of MH17 using open-source information.

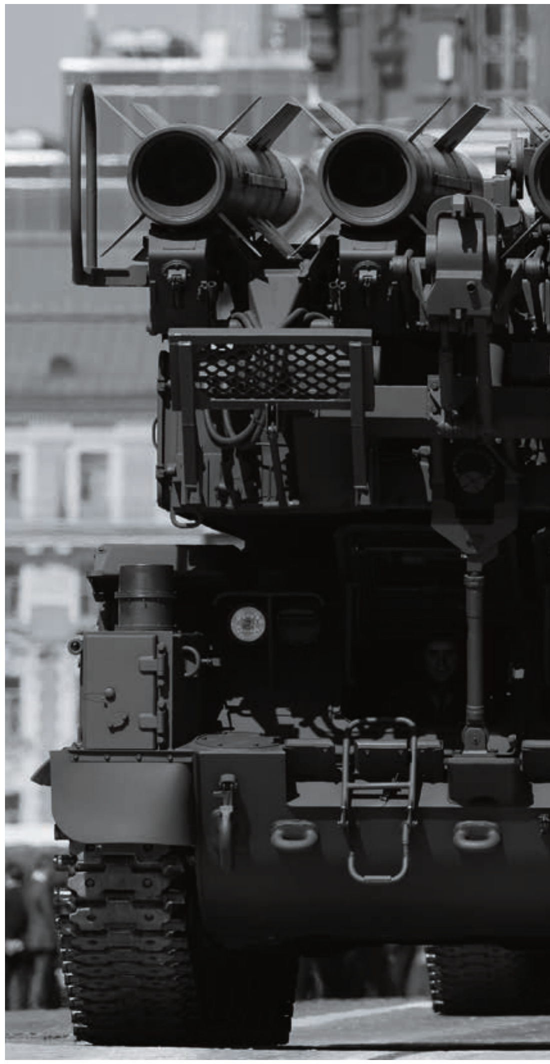

Weapons of mass deception

Russian Buk-M2 surface-to-air missile vehicles roll through Red Square during the Victory Day military parade in Moscow on May 9, 2015.

PIROSCHKA VAN DE WOUW/ GETTY

keyboard partisans in the war dominating Europe's largest country. After the annexation of Crimea, the subsequent invasion of eastern Ukraine, and the MH17 shoot-down, the world's scrutiny of Russian behavior in the region dwindled. Yet Ukraine—site of the continent's first military incursion by a neighbor since Stalin subjugated Eastern Europe behind the Iron Curtain—is, as U.S. Deputy Secretary of Defense Bob Work said in a speech in 2015, "an emerging laboratory for future 21st-century warfare."

The 34-ton Buk-M1 TELAR ("transporter erector launcher and radar"), and its bodyguard of irregular troops, rolled through the southeast corner of what the Yale historian Timothy Snyder has christened the Bloodlands. Here, in living memory—between 1933 and 1945—a hellish amalgam of Nazis and Soviets (sometimes collaborating, more often at war) were responsible for 14 million civilian deaths. "During the years that both Hitler and Stalin were in power," Snyder writes, "more people were killed in Ukraine than anywhere else in the Bloodlands, or in Europe, or in the world."

Half an hour after MH17 took off, another Ukrainian curator, @WowihaY, tweeted that the convoy had passed through his hometown of Torez, 45 miles east of Donetsk, headed to the city of Snizhne. There, the Buk was unloaded from a white Volvo low-loader truck before continuing south under its own power. Passing through checkpoints held by Russian-backed insurgents, it set up in a field and, at 4:20 P.M. local time, fired a

YURI KADOBNOV/GETTY

1,500-pound missile 33,000 feet into the air. Carrying a high-explosive fragmentation warhead weighing 154 pounds, it nearly reached Mach 3. On board MH17 were 15 crew and 283 passengers, including 80 children in 20 family groups and a party, led by the virologist and former International AIDS Society president Joep Lange, heading for the 20th International AIDS Conference in Melbourne.

The warhead exploded around four meters to the upper left of the airplane's nose. Dying bodies fell "like confetti" for around 90 seconds. One female victim crashed through the corrugated roof of a house into a kitchen. Autopsies found hundreds of metal fragments in the captain's corpse, another 120 in the first officer, and a bow-tie-shaped fragment—unique to the Buk-M1's 9N314M warhead—embedded in one of the flight crew.

For long days, governments scrambled to negotiate access with hostile irregular forces, probably composed of what Galeotti calls "a mix of regular Russian units, ad hoc collections of nationalists and adventurers, and everything in between." These auxiliaries, largely organized by the GRU (the Russian army's foreign military intelligence agency), now controlled what a spokesman for investigators with the Organization for Security and Cooperation in Europe (OSCE) called "the biggest crime scene in the world." The hot sun glittered. It was summer. It was very warm.

PENETRATING THE SMOG OF WAR

Three days earlier Eliot Higgins, a highly regarded citizen journalist, had launched his crowdfunded project Bellingcat in beta. Bellingcat would use open-source information, he promised, "to investigate, collaborate, and report on worldwide issues that are being underreported and ignored … Syria, Iraq, Turkey, Kurdistan, Nigeria, Jihadists, Shia armed groups, the U.K. phone hacking scandal, police corruption, and more." That "and more" swiftly became the downing of MH17. Bellingcat could have been designed for the challenge. Less than six hours after the shoot-down, Higgins had found, identified, archived, uploaded, and shared a 35-second video—titled in Russian "The Murder Weapon Malaysians Snizhne"—in which the Buk, now led by a single vehicle, rumbled ominously through Snizhne. Two years later, the Dutch-led international Joint Investigation Team (JIT) would use the video in its findings.

The official investigation by the Dutch Safety Board, conducted in parallel with the JIT's, embodied a century of hard-won knowledge about air accidents. Over 15 months, the $4.8 million investigation reconstructed substantial parts of the Boeing 777. A wealth of expertise fortifies every part of the 279-page report and its 26 appendices, showing precisely how flight MH17 was destroyed. In their effort to find out what happened—and who was responsible—the JIT's hundreds of investigators have, among other tasks, processed 1,448 pieces of wreck-

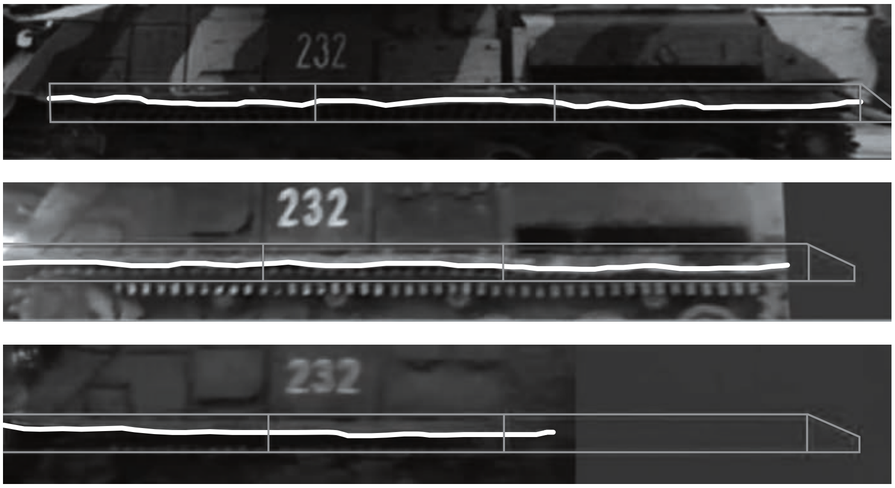

Side skirt fingerprint

Bellingcat researchers developed an innovative way to "fingerprint" vehicles as part of their successful effort to identify Buk 332, which shot down Flight MH17.

age, heard over 200 witnesses, analyzed 150,000 intercepted calls as well as half a million photos and videos, and produced over 6,000 reports. Although determined to keep their powder exceptionally dry for future criminal prosecutions, last September they presented preliminary results. After noting the efforts of "research collectives like Bellingcat," they reached an unequivocal conclusion: a Buk-M1 TELAR, armed with 9M38M1 missiles carrying 9N314M warheads, traveled from the Russian Federation into Ukraine, fired from a launch site roughly halfway between the villages of Pervomais'kyi (May Day) and Chervonyi Zhovten (Red October), and then returned to Russia.

On a shoestring budget, using social media and satellite photography, and tapping into a network of obsessives, Bellingcat's detective work has produced impressive results. In a series of reports, participants identified the actual Buk—unit designation number 332—and its battalion in Russia's 53rd Anti-Aircraft Missile Brigade. Comparing dozens of Buks, and analyzing photos shared on Vkontakte between 2009 and 2013, they homed in on seven characteristic markers. These included exhaust deposit patterns, dents, the arrangements of cable connections to the missile erector, fonts (and spacing) on the digits, and the vehicles' mix of hollow and spoked wheels on each side. A Bellingcat with an intelligence background developed an innovative type of "fingerprinting": using 3-D software "to solve the problem of comparing two vehicles with perspective-distorted photos," he noticed there were unique patterns of deformation in the protective rubber side skirts above the wheels.

Bellingcat was also the first to publicly describe the route the Buk took through Russia in late June and into and out of Ukraine before, during, and after July 17. The project has since identified several dozen soldiers associated with Unit 32406—the 53rd Brigade—by piecing together content and friend lists on Vkontakte, cross-referenced with posts on a forum for the often anxious mothers and wives of soldiers. (The murdered Russian journalist Anna Politkovskaya wrote movingly on this subject in her essay "My Country's Army, and Its Mothers.")

THE PENUMBRA OF UNCERTAINTY

None of this cuts much ice in Russia. The Kremlin's fog machine went into overdrive. The full panoply of Russian state media, troll farms, semi-automated botnets, and what Russian novelist Nikolai Leskov called "*useful fools* and silly enthusiasts" began their murky work. The Russian government's response to the shooting down of MH17 was a charade, wrapped in a travesty, inside a miasma: a relentless campaign of abuse and deceit, trying to entangle every fact of the matter in a net of disinformation. Numerous attempts were made to hack the Dutch Safety Board. Several Bellingcats experienced spear-phishing attacks. Other targets included French and U.K. TV channels, NATO, OSCE, and the Polish, Dutch, Finnish, and Norwegian governments, as well as German political parties.

The primary "threat actor" was a cyber-espionage group called Fancy Bear (which has several names, including Tsar Team, APT28, Strontium, and Iron Twilight): Russia-based, and in all probability controlled by the GRU. As during the

BELLINGCAT

operations against the U.S. election last year, Fancy Bear seemed careless about disguising its activities. (FBI director James Comey, testifying to the House Committee on Intelligence in March, called the group "very noisy.")

Fighting this cyber-espionage is CrowdStrike's Dmitri Alperovitch (in 2013, one of *MIT Technology Review*'s 35 Innovators under 35). He was the lead computer security consultant on the DNC hacks and has been instrumental in identifying major Chinese and Russian hacking groups. Alperovitch grew up in Russia until his family moved to the United States in 1995. Like many people of Russian origin, he has strong feelings about disinformation. "The power of cyber," he says, "isn't the 'cyber Pearl Harbor' scenario—which we've been talking about for 25 years now and hasn't happened. The real power is in information."

Alperovitch thinks Russia gets "the true nature of the battlefield" in a way the West does not: "They've been thinking about this for a very long time—it actually goes at least as far back as the Tsarist era in the 1860s, when they created one of the first modern intelligence agencies, the Okhranka." After the 1917 revolution, when the Bolsheviks opened the Okhranka's archives, "they were shocked to discover how infiltrated they were and how much disinformation had weakened their movement," he says. "They modeled the KGB on the successes of Okhranka. So they didn't invent it—they stole it."

The highest-ranking Soviet-bloc defector to the West—Lieutenant General Ion Mihai Pacepa, the former chief of Romania's espionage service—has exposed even deeper roots. In a book he recently coauthored, *Disinformation*, Pacepa cites the Marquis de Custine, Russia's Tocqueville, who wrote in 1839: "Everything is deception in Russia." Custine quotes a distinguished and well-traveled Russian diplomat quietly confessing, "Russian despotism not only pays little respect to ideas

"The Russians have, fortuitously for them, simply stumbled on a truth of the century."

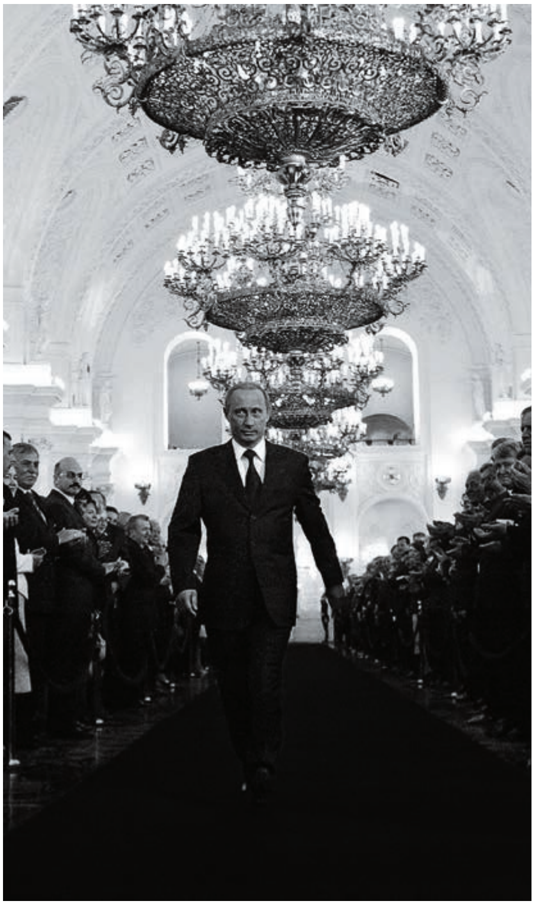

ALEXANDER ZEMLIANICHENKO/GETTY

The Great Khan

Vladimir Putin strides down St. George's Hall in the Kremlin's Grand Palace during his inauguration ceremony on May 7, 2004.

The lives of others

Wreckage from the downed Boeing 777 Malaysian airliner in a field in eastern Ukraine.

The culprit

Sergey Dubinsky, a veteran of Russia's wars in Afghanistan and Chechnya, was found to have played a key role in the missile launch.

and sentiments, it will also deny facts; it will struggle against evidence, and triumph in the struggle!" The tsar, and then Lenin, banned Custine's work.

The U.S. diplomat George Kennan, whose "Long Telegram" to the U.S. State Department shaped the Cold War and NATO, so admired Custine that he wrote a book about him in 1971. Kennan thought that much of Custine's analysis still rang true: "the neurotic relationship to the West; the frantic fear of foreign observation; the obsession with espionage; the secrecy; the systematic mystification; the general silence of intimidation; the preoccupation with appearances at the expense of reality; the systematic cultivation of falsehood as a weapon of policy; the tendency to rewrite the past." (Among Putin's methods, not least is the effort to ban his-

BELLINGCAT; ROB STOTHARD/GETTY

tory books as part of "the war of memories.") Kennan called Russians "one of the world's greatest peoples," but he retained a clear-eyed mistrust of their leaders.

To sunlight we can now add another powerful disinfectant: global, peer-to-peer, open-source investigation.

EXPECTATION IN THE AIR

"Answer in kind!" commands Edward Luttwak, the eminent if colorful military strategist and historian. Speaking from Moscow, he suggests that we respond aggressively to the global hacking of truth. "There are ample opportunities to hit back," he says, "but nobody is using them. There are a thousand stories here, openly circulated." (He shares one picked up in the Beluga caviar bar from two billionaires the previous night.) Putinism can be likened to the Golden Horde "in advanced Mongol form," he says. "It's not just the Great Khan who must have billions of dollars: now the companions of the Court must also be multibillionaires." He suggests using these stories "to 'disassemble' Putin."

To sunlight we can now add another powerful disinfectant: global, peer-to-peer, open-source investigation. On the day Bellingcat opened for business in 2014, Russia began an artillery bombardment from within its own borders, using its own equipment and soldiers. It still lies about the barrage, as it does about most of its actions in Ukraine. Last December, Bellingcat fired back a salvo of truth: an interactive map showing hundreds of artillery strikes on Ukraine made from Russian territory. Open-source investigation was also used by Russia's opposition leader Alexei Navalny in a recent, virally shared video that exposed the scale of Russian corruption. Bolstered by drone footage documenting the leadership's spoils, the facts brought out tens of thousands in protests across Russia.

Earlier this year, Bellingcat identified the man who called the instrument of MH17's destruction "my Buk-M." In several tapped phone calls released by SBU (Ukraine's secret service) and the JIT, he was called Sergey Petrovsky. But Bellingcat uncovered his real identity: Sergey Dubinsky, a veteran of Russia's wars in Afghanistan and Chechnya. The calls begin as MH17 passengers are embarking: Dubinsky can be heard coördinating several still-unidentified people as they head to the launch site. Calling himself "Bad"—from his call sign, "Bad Soldier"—Dubinsky would later prove his identity to a skeptical user on a forum by confirming he owned a black Peugeot 3008: Bellingcat found a dash-cam video showing a black Peugeot 3008 leading the missile-launcher convoy east of Donetsk. Bellingcat also found evidence of Dubinsky acknowledging that he was indeed the voice on the recordings. Instrumental in smashing apart the lives of others, Dubinsky responded to the Bellingcat reports with an e-mail to the BBC, sent from his home in Russia, that was openly contemptuous, describing his "Homeric laughter."

Today, at any one time, facts keep a city of several hundred thousand people safely in the air. Most of those air dwellers carry smartphones equipped with the GPS technology that Ronald Reagan accelerated into civilian use in response to Russia's shoot-down of KAL Flight 007. Perhaps we could gain something similar from MH17: a better global positioning system, this time for information. If it is to work, it is unlikely to be entirely technological. ✚

John Pollock has written for MIT Technology Review *about the role of social media in the Arab uprisings ("Streetbook," July/August 2011) and that of civilians in the Libyan information war ("People Power 2.0," March/April 2012).*

James Allison Has Unfinished Business with Cancer

The pioneer of immunotherapy asks:

Why do most patients fail to respond to the newest cures?

By Adam Piore

R. KIKUO JOHNSON

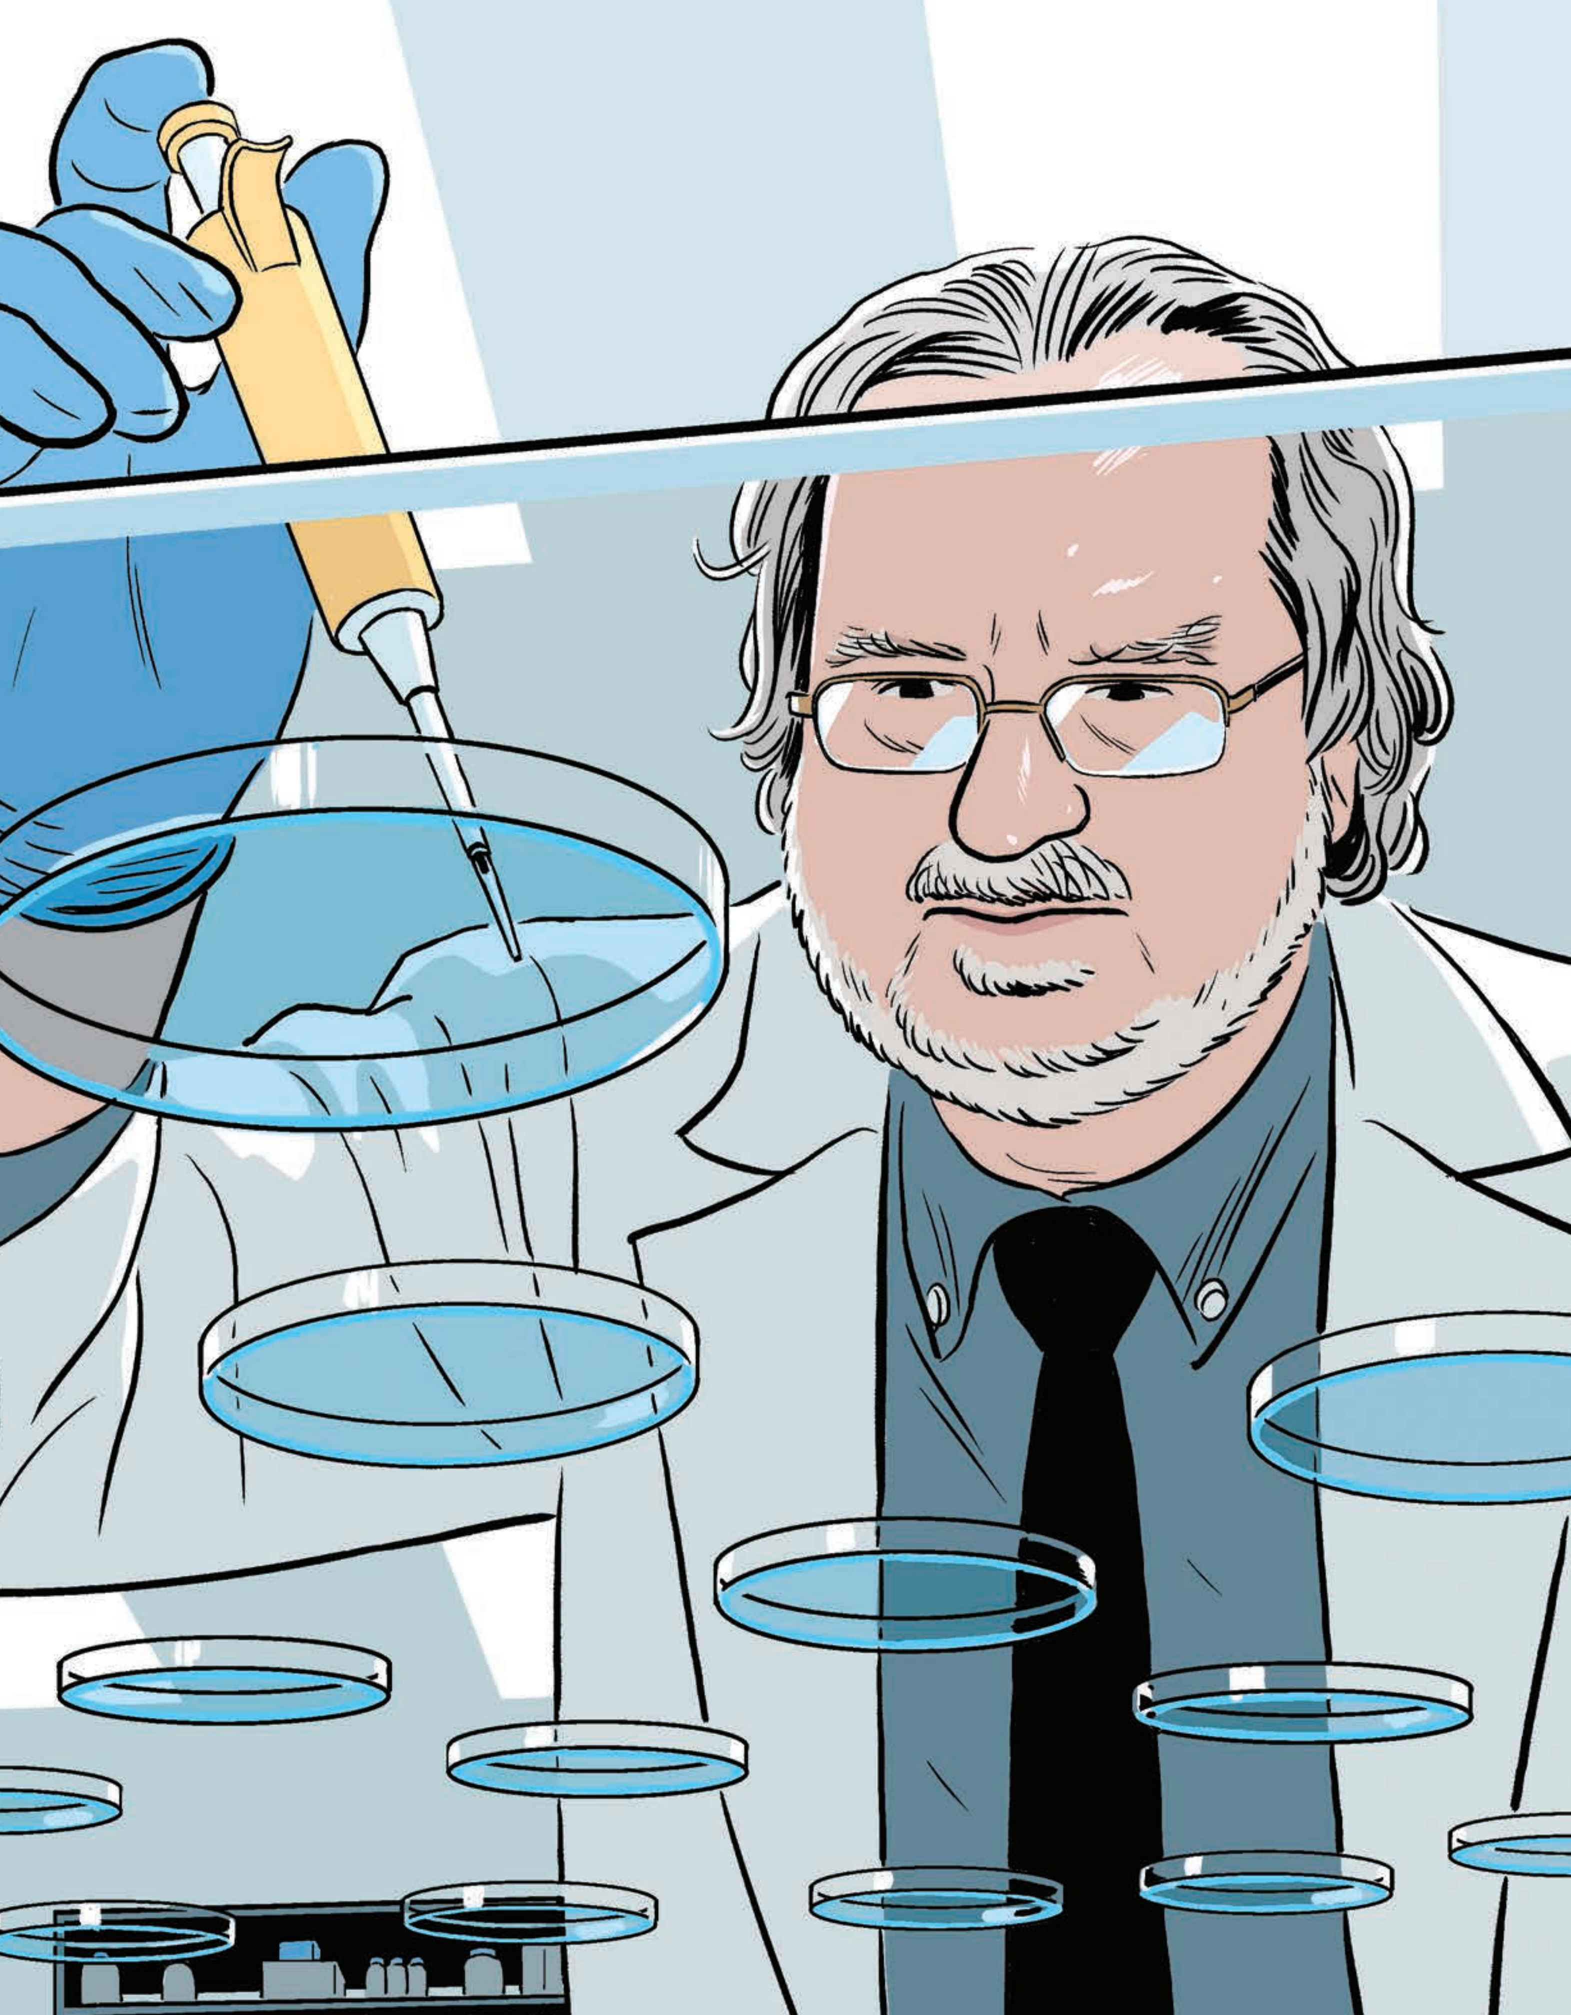

On the day I arrive at MD Anderson Cancer Center in Houston to meet James Allison and his longtime collaborator Padmanee Sharma, they are nowhere to be found. The previous day, one of their colleagues informs me, Allison was summoned up on stage by Willie Nelson, in front of 60,000 people at a rock festival in Austin, to deliver a harmonica solo. They are still on their way back.

By now, Allison is almost used to adulation. There are even murmurings that his work in cancer immunotherapy might win him the Nobel Prize. Twenty years ago, he was the first to show it's possible to turbocharge the body's response to cancer with a drug that releases the immune system so that it destroys tumors on its own.

The drug he identified to do that, called Yervoy, went on sale in 2011 to treat metastatic skin cancer. In lucky patients, it causes otherwise fatal tumors to melt away. By last year, worldwide sales of Yervoy and two newer drugs had reached $6 billion a year, and the medications had been given to more than 100,000 people. This transformative new class of immunotherapy agents, known as checkpoint inhibitors, is acknowledged to be the most important advance against cancer since chemotherapy.

Allison, who is 68, is an unimposing man, with a slight Texas drawl and a stringy mane of white hair. He still finds it hard not to cry when he meets cancer survivors saved by his discovery. But I had gone to talk to him about unfinished business. That is because for every miracle cure, for every Jimmy Carter or 22-year-old melanoma patient pulled back from death, there are many more people who, for reasons that no one understands, can't be saved. Of all patients dying from all types of cancer in America this year, only one in 12 would be expected to benefit from any immunotherapy drug. Some even argue that direct-to-consumer marketing, including a Super Bowl ad, has created dangerous expectations. Patients cashing in their last chance will, more likely than not, find themselves among the large majority for whom drugs like Allison's don't yet work.

Allison has known about the shortcomings longer than anyone. He says they dampen any sense of triumph and shadow him at award banquets. Sometimes, he stays awake at night. "About 22 percent of melanoma patients that get a single round of treatment with Yervoy are alive 10 years later," he said after receiving a Lasker Award in 2015, and then added solemnly: "We got to get that up, and we got to do it in more kinds of cancer."

At MD Anderson I was introduced to what Allison calls the "platform." It is a large-scale effort to determine why the immune system at times acts like the perfect weapon but in other cases fails to kick into action. Sharma, a Guyanese immigrant and practicing cancer doctor, oversees the collection of tumor samples from 100 of Anderson's 165 cancer trials that involve immunotherapy. The tissue is then scrutinized by her lab and Allison's for clues to how the battle is proceeding. "What is the immune response doing that leads to tumor rejection? What is the immune response doing that it stops rejecting the tumor and [it] starts growing again?" Sharma asks. "Those are big questions that we still need to understand."

The answers can't come too soon for some. The pharmaceutical industry and research institutions are in the midst of a pell-mell sprint into thousands of clinical trials based on new immunotherapy agents. As of October, by one tally, more than 166,736 patients were being sought to fill slots in studies of drugs involving a single protein, called PD-1. The overall number of immunotherapy trials probably tops 3,000, says Jeff Bluestone, an immunologist at the University of California, San Francisco, who also serves as president and CEO of the Parker Institute for Cancer Immunotherapy.

But a growing number of researchers fret that the flood of clinical trials is uncoördinated, redundant, and potentially counterproductive. That is because in many cases, the basic science remains little understood. "This is not sustainable," Ira Mellman, the keynote speaker at the annual meeting of the Society for

A Time Line of Cancer Treatment

Over 150 years, doctors learned to treat cancer with surgery, x-rays, chemotherapy, and vaccines. Immunotherapy is the latest weapon in the arsenal.

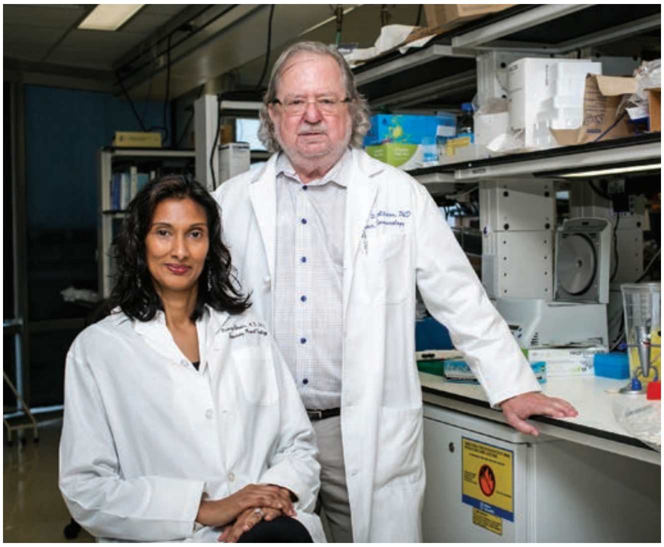

James Allison with his wife and scientific partner, Padmanee Sharma.

Immunotherapy of Cancer, told his colleagues when he took the stage last fall. Mellman, a vice president at the biotech behemoth Genentech, put up a byzantine diagram, consisting of concentric circles crammed with small type. The visually overwhelming slide showed trials under way to test immune-boosting therapeutics. His industry, he said, is "[throwing] plates of pasta against the wall, and hoping that something is going to stick."

Mellman told me that while Allison hadn't invented immunotherapy, his drug had been the one that clarified its potential. Now, he says, Allison's is one of the "few serious efforts" to better understand the mechanisms by which the immune system is killing cancers and the reasons why, too often, it is still not seeing them. "We would have a much better shot at doing what's best for patients, doing best for science, if we understand mechanisms," he told me. "You can just wildly try different things and hope that something works, or you can go back and try again and understand the basis of all of this. Until we know that, we're not going to really understand why some respond and some don't."

CHECKPOINT DISCOVERY

Cancer is personal for Allison. At 10, he held the hand of his mother, Constance, in tiny Alice, Texas, and wondered at the burn marks up and down her neck. He had not expected her to die. Only later did he learn that the marks were from radiation, and that cancer had killed her. By the time he was 15, cancer had consumed two of his uncles.

When Allison first began to chart a scientific career, he says, he recoiled from cancer. Back then, it seemed, there were few real clues. And immunology, the field he had picked, had a particular reputation for serving up fool's gold when it came to the disease. "I couldn't get any purchase on it," he recalls. "I wasn't going to go crashing into something until I knew how it worked."

At that time, in the 1970s, T cells—those tiny assassins that allow the body to fight off infections—had only recently been discovered. Allison was fascinated to learn there were molecular-level sentinels that patrolled the human body looking for trouble—that "if they see something wrong, they just deal with it." He thought: "What could be cooler than that?"

The existence of such immune cells did raise an obvious question: if T cells were designed to protect the body by killing infected and diseased cells, how was it that cancer managed to elude them? By then, there were hints that sometimes tumors did in fact succumb. In the 19th century surgeons had inoculated cancer patients with heat-killed bacteria, with inconsistent results. In 1980, a *Time* magazine cover spotlighted a scientific frenzy around a molecule called interferon, which sends the immune system into overdrive. But the treatment was indiscriminate, as likely to harm a person as heal. "It was crazy, because people were doing things and they didn't understand how they worked," Allison remembers. "People just said: 'Oh, well. It causes T cells to grow. So we put tons of it into people.'"

Allison instead began studying the molecular receptors present on the T cell's

1880–1957

c. 1880: Surgeon William Stewart Halsted argues that the recurrence of cancer after surgery is due to traces that remain. He helps pioneer the radical mastectomy.

1896: Emil Grubbe uses an x-ray tube to perform radiation therapy on Rose Lee, a mother of four with breast cancer.

1949: Mustard gas is approved by the FDA as the first chemotherapy after it is shown to destroy malignant white blood cells in lymphoma patients.

1957: The first bone marrow transplants are performed in Seattle. Although all six patients die within 100 days, the technique is a breakthrough.

ILANA PANICH-LINSMAN; R. KIKUO JOHNSON

“[WE WERE] JUST SOBBING, AND EVERYBODY WAS REALLY HAPPY.”

surface. One of his most important findings was to locate a receptor called CD28 that acts like a gas pedal. When it gets engaged, it is one of two key signals—in addition to a receptor that actually locks onto a tumor cell and functions somewhat like an ignition switch—that a T cell needs to initiate an attack.

Even when those switches were flipped to the “on” position, however, such attacks were often short-lived and sometimes failed to start up at all. By 1992, Allison thought there might be a third switch. The most likely candidate: CTLA-4, a mysterious receptor sometimes spotted on T cells. But both Allison and Bluestone, the UCSF immunologist, found that this molecule behaved unexpectedly. When proteins bound to it, it didn’t turn a T cell on—it turned it *off*. These molecular brakes were called checkpoints.

Scientists subsequently demonstrated why evolution might have favored checkpoints. When they created mice lacking CTLA-4, their T cells ended up attacking their own bodies after an infection. Without an off switch, the mice “died within a few weeks, of massive autoimmune disease,” Bluestone recalls.

Bluestone initially saw the chance to develop new types of immune-suppressing drugs—say, for organ transplant patients. But Allison saw a different possibility. Releasing these brakes might strengthen the immune system’s response against cancer. One of Allison’s graduate students had already developed an antibody able to stick to a T cell’s CTLA-4 receptors, essentially jamming the switch. Allison instructed a postdoc to inject the antibody into mice riddled with tumors. The results, he recalls, “were spectacular.”

“The tumors were cured,” Allison says. “I mean, it was 100 percent and zero percent—no statistics necessary.”

MIRACLE DRUGS

The drug, the first of the checkpoint inhibitors, would become known as ipilimumab or Yervoy, and it is now sold by Bristol-Myers Squibb, a pharmaceutical company headquartered in Manhattan. Human studies began around 2000 on 14 patients stricken with metastatic melanoma, who were steeling themselves for their finals days in hospice. But after the trial began, three saw their tumors shrink. Allison, who moved to New York City’s Memorial Sloan Kettering in 2004 to be closer to the trials, soon met one of the patients his drug had saved. Sharon Belvin was in her 20s, and had just finished college and gotten married, when metastatic melanomas appeared in her lungs, liver, and brain. She was terminal by the time her physician enrolled her in the first phase II clinical trial. The day Allison met her, she’d been in remission for a year.

“She hugged me,” Allison recalls. “Her husband hugged me, and her mother and father were there, and they all hugged me. It was just sobbing, and everybody was really happy. I walked to my office and I had a lot to think about. I cried all the way there.”

Allison says by that time he was aware of his drug’s limitations. It didn’t help everyone, and it didn’t work in most cancers. And if he needed a reminder of the stakes, it came in 2005, when Allison’s brother succumbed to prostate cancer after eight years. The same year, doctors found early-stage cancer in Allison’s own prostate. He had surgery rather than chance drug treatment.

As soon as cancer researchers learned that Yervoy worked on some previously incurable patients but not on others, many asked the obvious question: was it possible the body had more than one checkpoint? Another molecule, called PD-1, was quickly identified and successfully targeted with checkpoint inhibitors. Allison’s Yervoy was approved in 2011 by the U.S. Food and Drug Administration for patients with melanoma. Three years later the FDA approved Merck’s PD-1 inhibitor pembrolizumab (Keytruda) and a similar drug, also from Bristol-Myers Squibb, called nivolumab (Opdivo). One or both have since been approved to treat some types of lung cancer, kidney cancer, and Hodgkin’s lymphoma, creating the most important new class of cancer drugs in a century.

GATLING GUN

The day I arrived at MD Anderson to tour the platform, an Argentinean immunologist, Luis Vence, greeted me in a fluorescent-lit hallway. Our first stop was a lab where he swung open the door

COURTESY OF JAMES ALLISON; R. KIKUO JOHNSON

of a refrigerator-size machine to reveal 28 black canisters arrayed around a central hub, like the bundled cylinders of a Gatling gun. When cancerous samples come in, they are treated with fluorescent antibodies designed to stick to CTLA-4, PD-1, and other molecules on the surface of immune cells. The machine can then, in a few seconds, use a laser to scan all 10,000 or so cells from a biopsy, count them, and separate them by type. Vence compared it to sifting through multicolored ping-pong balls.

In a nearby lab, one of his colleagues, Jorge Blando, directed me to a microscope through which I could see a panorama of a cellular battle under way. The slide contained a slice of bone marrow riddled with tumors. These were recognizable by their larger, fuller-shaped cells. Among them were the tiny immune cells, stained brown, that had infiltrated and begun to attack. Others seemed to hover around the periphery. How many eventually make it in—and how long they survive to keep fighting—determines whether the tumor is defeated.

"What you are looking at in cancer is natural selection at a high speed," Vence says. "When you treat it with chemotherapy, maybe you destroy 99 percent of the tumor. But the 1 percent that is left is resistant to chemotherapy. That's the one that grows back and basically kills you." This explains why even the latest targeted drugs—those designed to hit very specific molecules on, say, a breast cancer cell—typically extend patients' life by only a few months.

Yet Vence and others believe that the immune system is inherently capable of spotting and countering any move a cancer makes. How else to explain how some advanced melanoma patients, who had tumors in their lungs or brain, are disease-free years after a course of Yervoy infusions? "The beauty of immunotherapy," Vence says, "is that the immune system can evolve at the same time, along with the tumor. It can keep up much easier."

It was Sharma who had the idea for the platform. When she began it, few volunteers were yet receiving Yervoy, then a relatively new and unproven drug. So Sharma persuaded patients slated to have less serious tumors removed by surgery to take small doses. A biopsy sample was collected before the drug was administered. Then, comparing the initial cancer and the excised tumor, the lab could use state-of-the-art technology to track the immune response and begin to examine why it didn't always work. Sharma's first finding came fairly quickly. In tissue from

1981–2006

1981: A vaccine against hepatitis B, which causes liver cancer, becomes the first cancer vaccine to reach market in the U.S.

1995: James Allison rids mice of tumors using a new type of treatment that unleashes the immune system: a checkpoint inhibitor.

1997: The antibody rituximab is approved to treat non-Hodgkin's lymphoma. It is the first molecularly targeted cancer drug.

2006: Cancer enters the genome era. Johns Hopkins scientists apply high-speed DNA sequencing to 22 tumors.

2006: Mass vaccinations begin against the human papillomavirus, the cause of cervical cancer.

In the late 1980s, James Allison began studying the molecular basis of T-cell behavior.

“I CAN’T BELIEVE WE MISSED THIS. THIS IS AMAZING ... I LOVE YOU!”

bladder cancers treated with CTLA-4 antibodies, readings from the Gatling gun showed that T cells possessing a molecule called ICOS were “off the charts.” Sharma’s reaction was elation mixed with confusion. T cells with ICOS on them had previously been found only in the tiny sacs in the lymph nodes known as follicles, and they were believed to suppress immune responses, not enhance them. Allison decided to engineer mice whose tumors triggered ICOS. In their tumors, CTLA-4 was four times as effective. ICOS, it turns out, was part of cascade that made T cells attack tumor cells more effectively.

“I can’t believe we missed this,” Allison remembers telling Sharma. “This is amazing.” He’d been scooped by his collaborator and felt blown away. They’d been spending more and time together, talking on the phone and working on science. Now he blurted out: “I love you!” Sharma recalls plowing forward with the conversation as if nothing had been said. But he had said it. The pair were married in a small ceremony in 2014.

With the help of the Boston venture capital firm Third Rock Ventures, they also started a company called Jounce Therapeutics that is developing a drug to increase ICOS levels. Human tests got under way last year, and although it’s too early to know how the drug is working, the idea has already been remunerative. Jounce went public in January, raising $117 million. Now Sharma drives a Tesla with a vanity plate that reads “ICOS.” On Allison’s Porsche, the plate says “CTLA4.”

ILANA PANICH-LINSMAN; R. KIKUO JOHNSON

At MD Anderson, where James Allison has his laboratory, researchers study why some people don’t benefit from immunotherapy.

2011–2016

2011: Ipilimumab, or Yervoy, is approved to treat advanced melanoma. It is the first checkpoint inhibitor to reach the market.

2015: Former president Jimmy Carter, at 91, has melanoma in his liver and brain. A checkpoint drug leaves him cancer-free.

2016: Recognizing “amazing advances” in immune therapy, President Barack Obama and Vice President Joe Biden announce a new “moonshot” to cure cancer.

A TIDAL WAVE

During the same meeting at which Mellman castigated the industry for spaghetti throwing, I saw Allison huddled over an iPad with another scientist, discussing some of the most recent findings he, Blando, and Sharma have made using their platform. They have been studying prostate cancer, in which no checkpoint drug yet seems to work. “What we found was that prostate cancer is almost a desert immunologically,” Allison says. “It’s a very cold tumor. There’s not much in there.” But Blando’s microscope has revealed that two drugs together might make the difference. Yervoy, he found, is necessary to drive T cells into the tumor, while the addition of a PD-1 drug makes sure they start killing. On the basis of these results and further research, Sharma and Allison convinced Bristol-Myers to combine the drugs in a clinical trial for advanced prostate cancer.

Many immunotherapy trials don’t have as much new preclinical research behind them. One reason is that drug companies are still exploiting the original checkpoint discoveries. Bristol-Myers’s Opdivo has been approved for eight different cancer “indications” in two years, which must be a record. “The pace of the clinical applications of the science is much faster than understanding mechanisms in the lab,” says Gregory B. Lesinski, a scientist at the Winship Cancer Institute of Emory University.

But racing ahead of the science can also incur huge penalties. Last summer, a test of Opdivo as a first-choice treatment for advanced lung cancer led to one of the greatest fiascoes in the company’s history. Bristol had organized a trial that, in seeking the largest market, had essentially taken all comers. Its competitor Merck chose to test its drug in only lung cancer patients whose biomarkers indicated they were most likely to respond. When Merck reported its results in June, they were so good that independent monitors said patients in a control group using chemotherapy could switch to the new drug immediately. Then, in August, Bristol acknowledged that its own test had failed to show a benefit. Shares of the company dropped by 20 percent, and Bristol’s research and development chief eventually stepped down.

The revival of immunotherapy now includes cancer-fighting viruses, genetically reprogrammed T cells, and vaccines designed to make tumors more visible to the immune system. Understanding the best way to put it all together is one of the crucial jobs ahead. At times, the explosion of new activity has tended to diminish the importance of Allison’s drug. Although it is still a billion-dollar-a-year blockbuster, Yervoy is now less often prescribed, in part because of side effects. One analyst called it the “iPod of immunotherapy”—a product overshadowed by the revolutionary change in thinking it caused. “Its importance would be hard to overstate in terms of what it has done to crystallize all the other activities,” says Mellman. “In my opinion, the idea that the immune system could target cancer didn’t start with Jim. But the field did.”

Once a year Allison packs a sold-out venue at the American Society for Cancer Research. There his own band, the Checkpoints, plays to doctors and scientists, nearly all of whom are converts to immunotherapy. Yet Allison still recalls the reviewer who, two decades ago, told a journal to reject his breakthrough paper because “we all know that immunotherapy’s crap. It’s never worked.”

Now that immunotherapy looks like the future, how far can it go? As I stood with Allison and Sharma in the MD Anderson parking lot, saying our good-byes, they seemed hopeful. Allison grabbed a piece of paper and sketched a graph. Start with everyone who has cancer, he said. Then, going out to the right, trace the survivors: how many are left after two months, six months, a year. It’s a line that, for most advanced cancers, drops relentlessly to the dust. But immunotherapy is lifting the curve. In melanoma, there are more and more long-term survivors. Allison calls it “raising the tail.”

“Ultimately, the goal is to try to get the survival rate as high as we can in as many different kinds of cancers as we can,” he says. Allison has finally gotten purchase on the monster that darkened his childhood. And he is not going to let go. ✚

Adam Piore is the author of The Body Builders: Inside the Science of the Engineered Human.

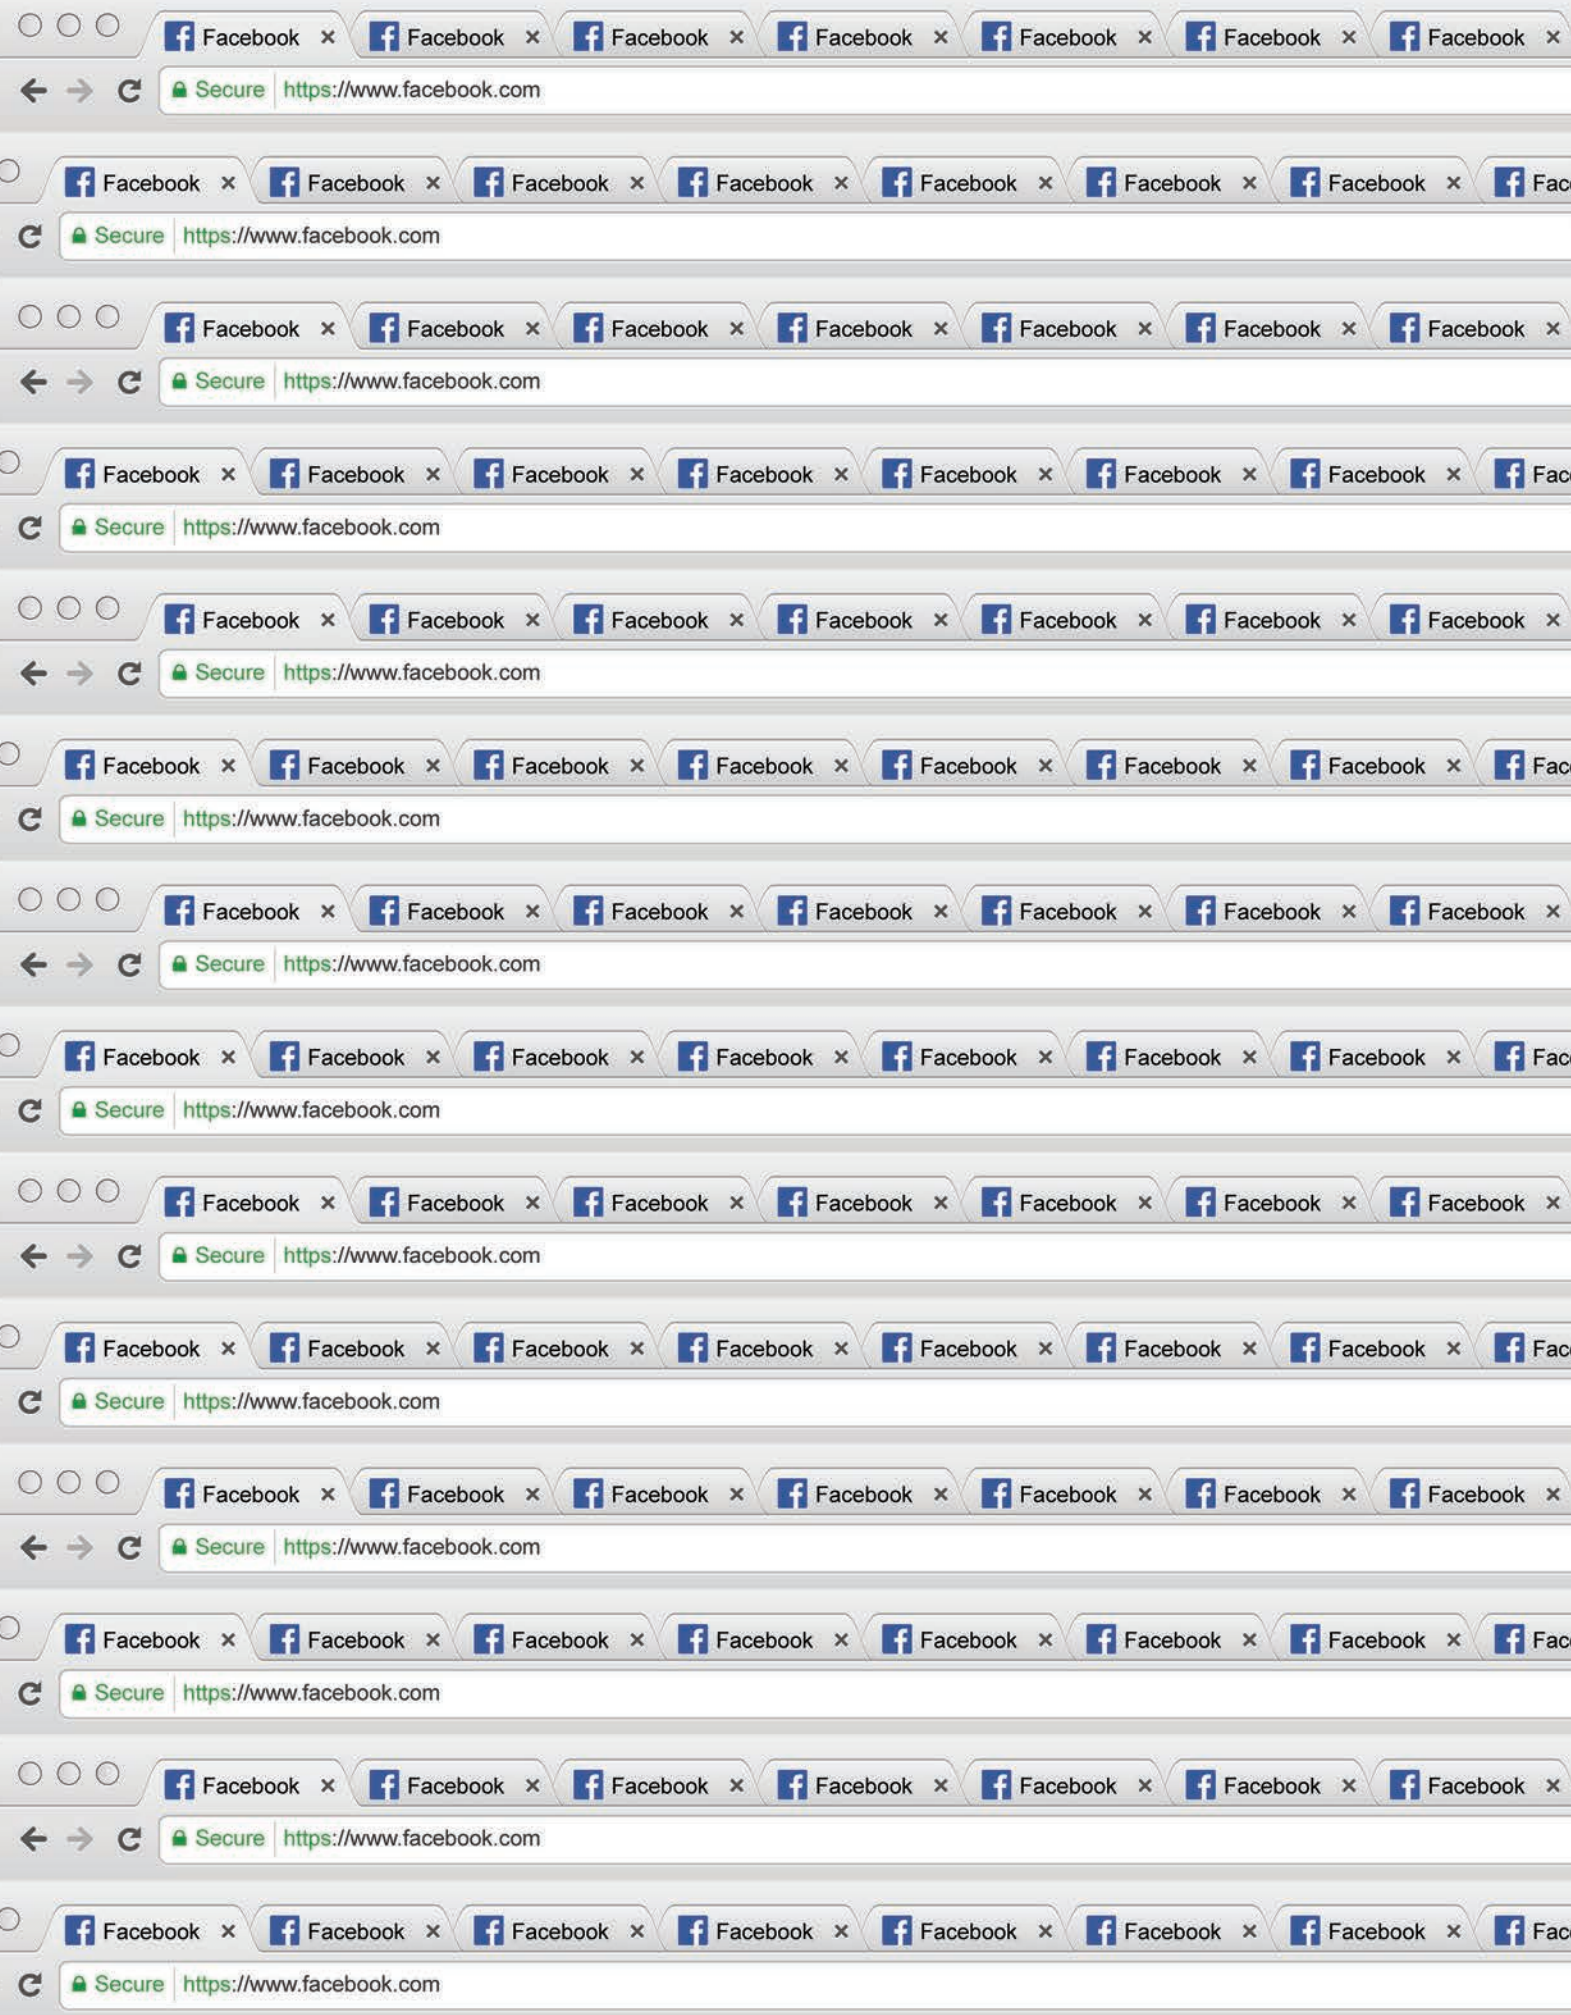

Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com
Facebook Facebook Facebook Facebook Facebook Facebook Facebook
Secure
https://www.facebook.com

Reviews

We Need More Alternatives to Facebook

Chastened by the negative effects of social media, Mark Zuckerberg says he will tweak his service and upgrade society in the process. Should any company be that powerful?

By Brian Bergstein

"They Have, Right Now, Another You"
By Sue Halpern
New York Review of Books
December 22, 2016

"Building Global Community"
By Mark Zuckerberg
February 16, 2017

About 10 years after TVs began to be ubiquitous in American homes, television broadcasting was a staggering financial success. As the head of the Federal Communications Commission observed in a 1961 speech to broadcast executives, the industry's revenue, more than $1 billion a year, was rising 9 percent annually, even in a recession. The problem, the FCC chairman told the group, was the way the business was making money: not by serving the public interest above all but by airing a lot of dumb shows and "cajoling and offending" commercials. "When television is bad, nothing is worse," he said.

That speech would become known for the pejorative that the FCC chairman, Newton Minow, used to describe TV: he called it "a vast wasteland." It's a great line, but there are other reasons to revisit the speech now, about 10 years after the emergence of another communications service—Facebook—that has become ubiquitous in American homes, a staggering financial success, and a transmitter of a lot of pernicious schlock. What's striking today is why Minow said the vast-wasteland problem mattered—and what he wanted to do about it.

As for why it mattered, Minow told the TV executives:

> Your industry possesses the most powerful voice in America. It has an inescapable duty to make that voice ring with intelligence and with leadership. In a few years, this exciting industry has grown from a novelty to an instrument of overwhelming impact on the American people. It should be making ready for the kind of leadership that newspapers and magazines assumed years ago, to make our people aware of their world.

On that point in particular, Mark Zuckerberg apparently would agree. "Are we building the world we all want?" he wrote in February, in a 5,700-word manifesto that reflected on the sometimes dubious role Facebook has been playing in civic life. Referring to its propensity to turbocharge hoaxes and to the way it tends to make news feel sensational, he wrote that Facebook's goal "must be to help people see a more complete picture" of the world.

But how to make a mass communication medium better for us? In 1961, Minow had a clear answer: "I believe that most of television's problems stem from lack of competition." He said he looked forward to seeing more channels becoming available through new technologies, such as UHF frequencies, pay TV, and international broadcasts. And he said he would look for ways to strengthen local stations that could best serve local communities. "I am deeply concerned with concentration of power in the hands of the networks," Minow said.

That's where Mark Zuckerberg would probably get a little uncomfortable. Because Facebook is all about concentrating power in one network—his, which he calls "a global community." If in reality Facebook tends to promote polariza-

MR. TECH

tion and tribalism, Zuckerberg seems to believe that can be fixed with a few tweaks. In his February letter he said Facebook would try to reduce sensationalism on the site and take other steps to help make people better informed and more engaged in democracy.

Zuckerberg doubtless means well, but the problem is not that we need a slightly better Facebook. It's that Facebook—a company worth $400 billion because it vacuums up information about our tastes, our shopping habits, our political beliefs, and just about anything else you might think of—is too powerful in the first place. What we need is to spend less time on Facebook.

Mesmerized

In his February letter, Zuckerberg essentially acknowledged what was obvious to anyone who had a Facebook account during the 2016 election: the social network has not exactly enhanced our democracy. The News Feed, the main scroll of posts that you see when you open Facebook, fueled hoaxes (which were overwhelmingly "tilted in favor" of Donald Trump, according to an analysis by Hunt Allcott of New York University and Matthew Gentzkow at Stanford), and it overfed people stories and memes that fit preconceived notions. On social media, "resonant messages get amplified many times," Zuckerberg wrote. "This rewards simplicity and discourages nuance. At its best, this focuses messages and exposes people to different ideas. At its worst, it oversimplifies important topics and pushes us towards extremes."

To try to counteract the fake-news problem, Facebook is now flagging hoax stories that are shared on the site with a warning that third-party fact checkers have declared them to be false. And in hopes of promulgating fewer stories that are apparently true but nonetheless uninformative, the company has adjusted the News Feed to give more weight to stories that people share after reading (or at least opening) them, rather than the ones they share after only seeing the headlines. The thinking is that a story shared largely based on the headline alone is less likely to be what Zuckerberg calls "good in-depth content."

Good for Facebook for trying these strategies. They fit with other civic-minded steps the company has taken in the past, such as encouraging people to vote and urging them to donate to the victims of floods and earthquakes. But the latest efforts probably won't do much to help create what Zuckerberg calls a more "informed community." The structure of Facebook works against that.

Facebook is fundamentally not a network of ideas. It's a network of people. And though it has two billion active users every month, you can't just start trading insights with all of them. As Facebook advises, your Facebook friends are generally people you already know in real life. That makes it more likely, not less, to stimulate homogeneity of thought. You can encounter strangers if you join groups that interest you, but those people's posts are not necessarily going to get much airtime in your News Feed. The News Feed is engineered to show you things you probably will want to click on. It exists to keep you happy to be on Facebook and coming back many times a day, which by its nature means it is going to favor emotional and sensational stories.

Zuckerberg's latest ideas probably won't create a more "informed community." The structure of Facebook works against that.

Why else would Facebook be increasing the prominence of video? In fact, one of its executives has suggested that within a few years the News Feed could be "all video." Surely some of the videos you'll see on Facebook will be in-depth documentaries, live feeds from news events, and other substantive material. But in general, showing us much more video from around the Internet does not feel like a way to promote more reasoned discourse.

As Zuckerberg himself noted in his February letter, most of what people come to Facebook for is ultimately social—"friends sharing jokes and families staying in touch across cities," or people finding support groups for everything from parenting to coping with a disease. For Facebook to be all that as well as a modern-day agora, a place of enlightened civic and political engagement, seems like a mismatch.

If you need a reminder that Facebook's primary reason for existence is not to enlighten you, consider the fact that the company catalogues a huge amount of information about you.

The behavior is not surprising—Zuckerberg claimed years ago that privacy was no longer a social norm—but the scale still astonishes. Last summer the *Washington Post* listed 98 of the data points that Facebook captures about its users. For example, by cross-referencing your behavior on Facebook with files maintained by third-party data brokers, the company gathers data on your income, your net worth, your home's value, your lines of credit, whether you have donated to charity, whether you listen to the radio, and whether you buy over-the-counter allergy medicine. It does this so that it can give companies an unprecedented ability to post ads that are presumably likelier to appeal to you. (I asked Facebook whether anything has changed to make the *Post*'s report no longer accurate; the company had no comment.)

This system may or may not work for advertisers, but it works very well for Facebook, which chalked up a net income of $10 billion on $28 billion in revenue last year. Does it work well for us? As Sue

NEIL DONNELLY

Halpern wrote in the *New York Review of Books*, the services that we get from Facebook are requiring us to give up something that is very hard to ever get back:

> Many of us have been concerned about digital overreach by our governments, especially after the Snowden revelations. But the consumerist impulse that feeds the promiscuous divulgence of personal information similarly threatens our rights as individuals and our collective welfare. Indeed, it may be more threatening, as we mindlessly trade ninety-eight degrees of freedom for a bunch of stuff we have been mesmerized into thinking costs us nothing.

When you look at Facebook that way, it's hard to root for the company to find ways to be a platform for *more* civic engagement. In fact, unless we think people should be required to shoulder whatever privacy costs Facebook decides to impose, it probably should not be the main place we go to find groups that, in Zuckerberg's words, "support our personal, emotional, and spiritual needs." Ideally, people would be able to form robust online communities and engage in the public square without letting any single company build a comprehensive dossier on them.

Lots of niches

What if we followed Minow's reasoning with TV in 1961 and decided that we ought to have many more powerful networks for disseminating ideas and shaping public discussions?

The first step would be to acknowledge that even with the seemingly limitless competition that already exists on the Internet, Facebook has an outsize role in our society. Sixty-eight percent of all American adults use it, according to the Pew Research Center. That compares with 28 percent for Instagram (also owned by Facebook), 26 percent for Pin-

terest, 25 percent for LinkedIn, and 21 percent for Twitter. And none of these other sites aspire to be as many things to as many people as Facebook does.

One of the interesting things about Minow's "vast wasteland" speech is that his encouragement of more competition helped inspire the expansion of public broadcasting in the United States. And perhaps it's time for similar efforts today, to support more varieties of social media.

These noncommercial alternatives would not have to be funded by the government (which is fortunate, given that government funding for public media such as PBS is in doubt these days). Ralph Engelman, a media historian at Long Island University who wrote *Public Radio and Television in America: A Political History*, points out that the creation of public broadcasting was led by—and partially funded by—prominent nonprofit groups such as the Ford and Carnegie Foundations. In the past few years, several nonprofit journalism outlets such as ProPublica have sprung up; perhaps now their backers and other foundations could do more to ensure the existence of more avenues for such work to be read and shared.

High-minded alternatives to Facebook have been introduced before. A now-defunct discussion site called Gather once got investment from American Public Media, a producer of public-radio programs. Among the platforms that still exist, Diaspora gives people ways to socialize without relinquishing control of their data. Parlio, now owned by Quora, was cofounded by a leading figure from the Arab Spring in Egypt to promote online discussions with "thoughtfulness, civility, and diversity." But we still could use more options that collectively counteract Facebook's enormous reach and influence and bring out more of social media's most constructive qualities—the way it connects us to far-flung people, information, and ideas.

Because noncommercial alternatives would be free of the imperative to capture as much information about your interests as possible, they'd be likelier to experiment with new ways of stimulating interactions between people. Maybe they would do away with the News Feed model that rewards virality more than importance. Perhaps some would be more reliant on algorithms to serve up stories and ideas, while others would rely on human curators to elevate discussion and eliminate abuse by booting trolls or deleting hoaxes.

Competitors to Facebook that harnessed the powers of social media only in an effort to make us wiser would probably be niche services, like National Public Radio and PBS. "Most people aren't that fussy," says Jack Mitchell, a journalism professor at the University of Wisconsin and the author of *Listener Supported: The Culture and History of Public Radio*. "PBS's market share is not that high. Public radio is a little higher. It's a minority taste."

As Facebook becomes more powerful than ever, other options are possible, and vital.

But having many more niche alternatives to Facebook could be exactly what we need. Even if none stole a significant chunk of Facebook's users, it might be enough to remind people that even as Facebook becomes more powerful than ever—rolling up massive profits and preparing to beam down Internet access to offline corners of the globe—other options are possible, and vital.

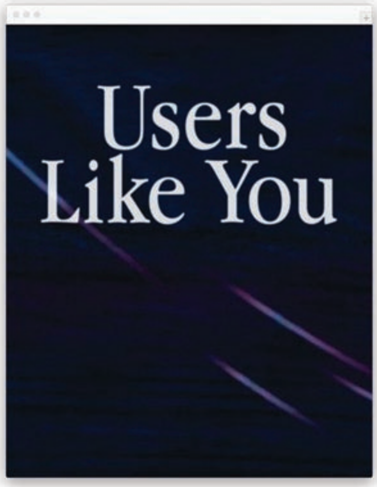

NEIL DONNELLY

Why are we finally now in what's often called a golden age of television, with culturally influential, sophisticated shows that don't insult our intelligence? It's not because broadcasters stopped airing schlock. It's because the audience is more fragmented than ever—thanks to the rise of public broadcasting and cable TV and streaming services and many other challenges to big networks. It required a flourishing of choices rather than a reliance on those huge networks to become better versions of themselves. As Zuckerberg wrote in February, "History has had many moments like today."

Brian Bergstein is editor at large for MIT Technology Review.

Me and My Troll

Years of unhappy interactions with an online commenter compelled the publisher of *MIT Technology Review* to rethink how his site hosts conversations.

By Jason Pontin

I have a troll. Writing as @zdzisiekm, or "Gus," or under other names, he has commented on stories at TechnologyReview.com 6,386 times and counting as of April 2017. As trolls go, he is unfailingly polite, and he doesn't violate our site's terms of service. Instead, he is reflexively, tendentiously wrong about a single topic, again and again. Gus is angry about our reporting on global warming and renewable energy technologies. His objections are notionally scientific, but they have a strongly ideological flavor.

Four years ago, commenting on "Climate Change: The Moral Choices" (see May/June 2013), @zdzisiekm characteristically wrote, "Having studied the relevant science literature quite extensively and in depth—I read hundreds of papers on the subject—there is no real 'climate change threat.' It's all trumped up—the actual published peer-reviewed science is clear on this … This is because in some countries [economists] are so keen to switch the economy away from fossil fuels, they'll go with any lie …"

Over our long association, Gus hasn't changed. Last January, after reading "What's at Stake as Trump Takes Aim at Clean Energy Research," he remarked, "None of the solutions fostered by the American Left a.k.a. the Democrats are affordable, safe, or … reliable. Adding intermittent energy sources to the grid has one effect only: it increases the cost of energy … As to safety, ask millions of bats and birds killed, blinded and fried in flight by windmills and solar installations. Ask people inconvenienced by the incessant annoying noise made by the windmills. Neither have these technologies created jobs … other than in China."

It's personal for @zdzisiekm; our interactions feel intimate and overheated. He has often denigrated my judgment and disparaged my qualifications. "This is really not your field," he recently wrote me in an e-mail.

I know who Gus is, because I tracked him down. We ask readers to provide some personal information before they can comment, and he wasn't hard to find. My troll is a sixtysomething technical advisor to the IT department of a large public university in the Midwest. He has not one but two PhDs—in electrical engineering and physics. He writes good research about computer architecture and bad poetry about cats. (I agreed not to use @zdzisiekm's real name for this story. "I know you know who I am," he said, "but I cherish my anonymity, and I don't want people to throw bricks at my window or dent my car.")

TOMI UM

When I asked Gus why he wastes so much time and spirit commenting on our site, he replied, "It doesn't take much of my time at all. I've got a personal database that I can quickly search for specific articles on various subjects of which I have, by now, tens of thousands." This is true. Like many trolls, @zdzisiekm cuts and pastes the same memes into many comments. He especially likes a post that begins, "All global warming seen since 1880 has been less than the natural centennial global temperature variability," followed by a cherry-picked list of papers from obscure journals with little or no peer review, meant to leave the impression that there is scientific debate about the causes of climate change.

Quizzed about his motives, Gus answered: "These are contentious and partisan issues. Let's not kid ourselves that they are not. This is precisely why I would expect balance in reporting on these topics, especially of *TR*. I suggested in the past that it may be a good idea to publish opposite views, side by side, as *WSJ* does sometimes. If *TR* did so, why, there'd be less reason for me to comment. In contrast, *TR* has been rather biased in its climate and energy articles." I tried to explain that we can't publish the "opposite views"—that climate is not affected by industrial emissions, and that if global warming did turn out to be real, humans could effectively respond when it became a problem—because those views are not true. To no avail: @zdzisiekm is a hard scientist; I am an ignorant editor.

We receive comments similar in their scorn, if different in politics, from readers who believe we publish the "PROPAGANDA AND LIES" of Monsanto and other creators of genetically modified organisms, or who are convinced we suppress the truth about the "filthy and unsound practice of vaccination" and its links to autism. What all such readers share is a conspiracist point of view: they think the scientific or economic consensus is in some way a hoax; that journal-

ists and academics are gatekeepers who enforce a dangerous orthodoxy, often for personal gain or party benefit; and that honest commenters must demonstrate that the Opposition cannot be silenced. Not all commenters on TechnologyReview.com are like this, but in recent years those who are have become more aggrieved, and they have discouraged other readers from commenting.

Our unhappy experiences with comments are common to most publishers. During the U.S. elections of 2016, when commenters were especially intemperate (either sincerely so, or because they had been paid to post, or were not humans at all but bots), the problem grew acute. Comment sections are now the digital spaces publishers have ceded to trolls, cranks, and conspiracy theorists of all kinds. Why do commenters do it? In "Conspiracy Theories," the legal scholars Cass Sunstein and Adrian Vermeule attribute conspiracist thinking to feelings of impotence: such theories are "especially likely to appeal to people who are cynical about politics, who have lower self-esteem, and who are generally defiant of authority." Commenting makes them feel less powerless and irritable. But why should a publisher put up with Gus or any troll? Why indulge them? What's in it for *me*?

TechnologyReview.com

The Coral Project

Perspective

Civil Comments

"Conspiracy Theories"
By Cass Sunstein and Adrian Vermeule
Journal of Political Philosophy
2008

"On Bullshit"
By Harry G. Frankfurt
Raritan Quarterly Review,
Fall 1986

Monomanical and grumpy impulses

While conceding that individual comments mostly have little value, defenders of commenting adduce three benefits to the activity. They argue that comments are the digital homologue of letters to the editor, and can be of intrinsic interest; that they are a way of "listening to your users," providing vital feedback about what are, in the end, products; and that comments serve business interests by goosing various measurements of reader engagement such as time on site or return visits, which in turn improve the performance of ads or the likelihood of selling subscriptions or memberships. In reality, the fraction of publishers' audiences who comment is so small and unrepresentative that only the first argument is valid, and then only on moderated sites with more or less knowledgeable readers, responding to quality journalism and information.

The reasons why publishers turn off comments are telling: since 2014, Vice, Recode, Reuters, *Popular Science*, *The Week*, Mic, The Verge, *USA Today*'s FTW, and many other sites have shuttered comments because they were too much work for little return. When NPR.org disabled commenting last August, the managing editor, Scott Montgomery, provided a quantitative rationale: "Far less than 1% of [a monthly audience of 25 to 35 million unique visitors] is commenting, and the number of regular comment participants is even smaller. Only 2,600 people have posted at least one comment in each of the last three months—0.003% of the 79.8 million NPR.org users who visited the site during that period." The ratio of commenting to reading on TechnologyReview.com is similar to National Public Radio's: in 2016, about 3,000 people commented on stories, out of 21,205,603 users of the site, making up just 0.014 percent of our total traffic. More anecdotally, those who did comment were more like @zdzisiekm than our larger audience: older, more monomaniacal, and grumpier.

In short, commenters aren't representative, and they're not numerous enough to meaningfully improve engagement. Worse, their comments demand constant pruning or deletion by dedicated staff or companies that specialize in beating back trolls, lest publishers acquiesce to nonsense or worse. Jonathan Smith, Vice.com's editor in chief, was more blunt than the civic-minded Montgomery when he explained why he was done with comments: "We don't have the time or desire to continue monitoring that crap moving forward."

Those sites that remain committed to comments have generally followed a limited number of strategies. Smaller publishers that disabled commenting on their own sites are reconciled to the fact that discussions moved to Facebook, Twitter, and Instagram. Recode's Kara Swisher said, "Things have changed; you have to change with them. Social media is just a better place to engage a smart audience that's not trolling. We got into a lot of trouble in our comments on different stories—attacks on our writers, just stupid things; it wasn't smart." Comments in social media are sometimes more civil, because many people use their real identities, which discourages trollish impulses. Larger publishers that choose to preserve on-site comments, including the *New York Times*, the *Guardian*, and the *Washington Post*, often constrain the problem by limiting either the number of stories with comments, the amount of time readers have to comment, or both. For instance, only 10 percent of stories on NYTimes.com have comments, and commenting is typically closed after 24 hours. Limiting the number of comments makes it possible for moderators to approve, reject, promote, or demote the best or the worst.

Technologists, as they will, have offered technological solutions to the problem of comments. The Coral Project, a collaboration between the *Washington Post*, the *New York Times,* and the Mozilla and Knight Foundations, provides open-source tools for newsrooms that want to build better commenting systems, including "Ask" and "Talk" functions. Perspective, created by Google's Counter-Abuse Technology Team and Jigsaw, a technology incubator at Alphabet that addresses challenges to free speech, uses machine learning to score how much any comment might tend to degrade or enhance a conversation. Civil Comments forces communities to rate a comment before it is posted. Finally, in an interesting experiment, NRKbeta, the technology site of the Norwegian public broadcaster, requires would-be commenters to prove they have understood a story by answering three multiple-choice questions before they can comment.

As for us, last year I grew so wildly dispirited at how *MIT Technology Review*'s stories had become part of America's endless, arid culture wars (and so frustrated with @zdzisiekm and a half-dozen other commenters) that we disabled commenting for four months in order to reimagine how we could host more enlightened and enlightening conversations. We, too, accepted that the most active commentary on our stories now occurred in social media, but we felt there was still a role for on-site comments. (Indeed, the two platforms can cross-fertilize each other in fruitful ways.) We believed that good comments could adorn and improve our journalism. But we suffered no illusions that commenters were representative of our broader readership or that comments served any direct business purpose. Building on Disqus and the Ask function in the Coral Project, our new strategy borrows widely from the solutions described above, and it is still a work in progress.

Illusionless strategies

We decided, in imitation of the *New York Times*, that readers would comment on only a few stories and then only for a while. Stories that might repay good commentary, such as our major features, essays, and reviews, would have comments, but those that might inflame partisan wrangling would not. We would choose to think of comments, whenever possible, as integral to the story: we wondered if we could construct whole stories around comments, or seed a conversation by inviting our smartest, most informed sources to comment. No one was doing this precisely, but some of the expert commentary at Ars Technica and The Information inspired us. We wanted readers to vote comments up and down, as readers once did in Gawker's Kinja. We knew that writers, Web producers, and the social media and community editor would have to be heavily involved in curating the comments; like the *Economist*, we wouldn't launch a thread and walk away.

"We don't have the time ... to continue monitoring that crap moving forward."

Finally, and most controversially, we decided that we wouldn't hesitate to censor comments or ban readers if they debased the site. That is, even if comments were politely expressed and relevant, and otherwise met our commenting guidelines, we felt we should be free to suppress their authors if they trolled us, posted bullshit, hijacked a thread, or contradicted known evidence. Screw @zdzisiekm and his gang, unless they behave as heirs to the tradition of civilized commentary. There is no inherent right to comment unless readers conform to various duties and responsibilities.

How is our commenting strategy working? Gus has responded well to the new regime, although his mind remains unchanged. He still comments nearly every day, but he says, "On my side, I've learned to comment with more precision and less, let's say, personal involvement." He argues less aggressively and more honestly, and he cuts and pastes less and links to defensible research more. Recently, he thanked me "for being reasonable about the whole business of commenting." We even have a bet: "If global temps drop all by themselves by 2030," he says, "you owe me a dinner at a restaurant of my choice; otherwise I owe you one."

Readers aren't universally happy, of course. When were they ever? Not long ago, responding to a story about an important project to create a "subcritical facility" to test small, transportable molten-salt-cooled nuclear reactors (see "MIT's Nuclear Lab Has an Unusual Plan to Jump-Start Advanced-Reactor Research"), "breister," one of @zdzisiekm's online pals, wrote, "Ah finally an article which did not disable comments. Censorship at its finest, complements [sic] of TR and their policy of squelching dissenting views."

You can't please everyone.

Jason Pontin is the editor in chief and publisher of MIT Technology Review.

Demo

The Inelegant Beginnings of a Dazzling New Wireless Technology

Within a few years, 5G networks could turbocharge your smartphone. But can they pass the tests cooked up by engineers at an office park in New Jersey?

By Elizabeth Woyke

Photographs by Lauren Lancaster

Live-streaming a virtual-reality broadcast. Downloading a 90-minute high-definition TV show to your smartphone in less than three seconds. Sending instant updates on road conditions to self-driving vehicles. These scenarios are impossible or prohibitively expensive on current cellular networks, but they should be feasible with the next generation of wireless connectivity, 5G. It promises to be 10 to 20 times faster than today's cell-phone networks.

Cell towers like these let AT&T engineers simulate how 5G will behave in the real world. Commercial deployment is expected to begin in 2020, after companies around the world hammer out technical standards.

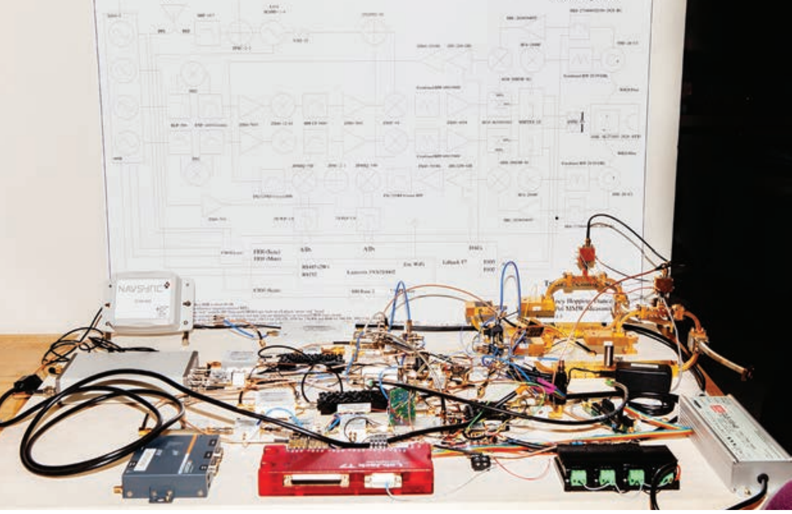

Research engineers hacked together these testing devices as part of their effort to see how 5G networks will perform in wooded areas and in inclement weather.

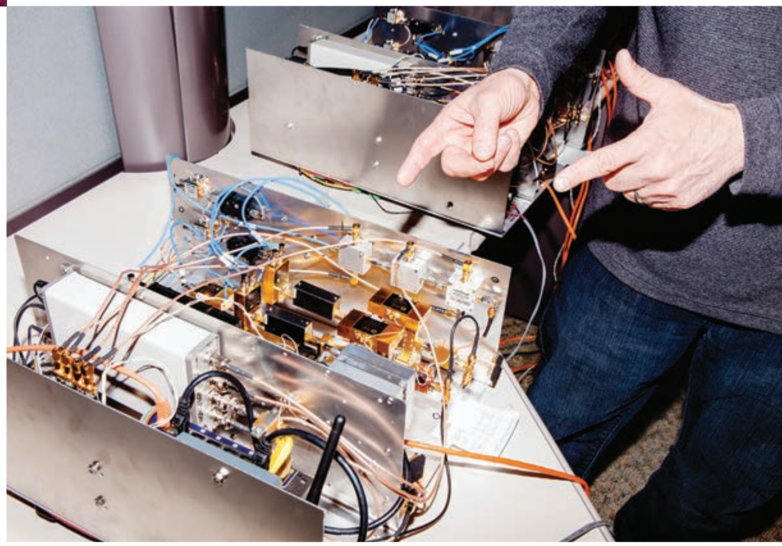

AT&T redesigned its measurement equipment to include 30 to 40 percent fewer components. Some parts were custom-built. Some were produced in-house.

That's because 5G will operate in a high-frequency portion of the radio spectrum, known as millimeter wave. It has a lot of available bandwidth and should make it possible for wireless devices to process data with minimal delays. But since its wavelengths are much shorter, it is more easily obstructed. And because it has never been used for consumer mobile services, carriers are still learning how 5G signals will behave in different types of terrain and weather. "We need to look at how the signals are affected by things like snow, rain, sleet, hail, maple trees, oak trees, and spruce trees, because each of those will be different," says AT&T research engineer Bob Bennett.

The problem: most 5G measurement equipment is so expensive, fragile, and

The roof of the main building on the Middletown campus is home to five 5G measurement systems, as well as other radios, weather stations, and solar panels.

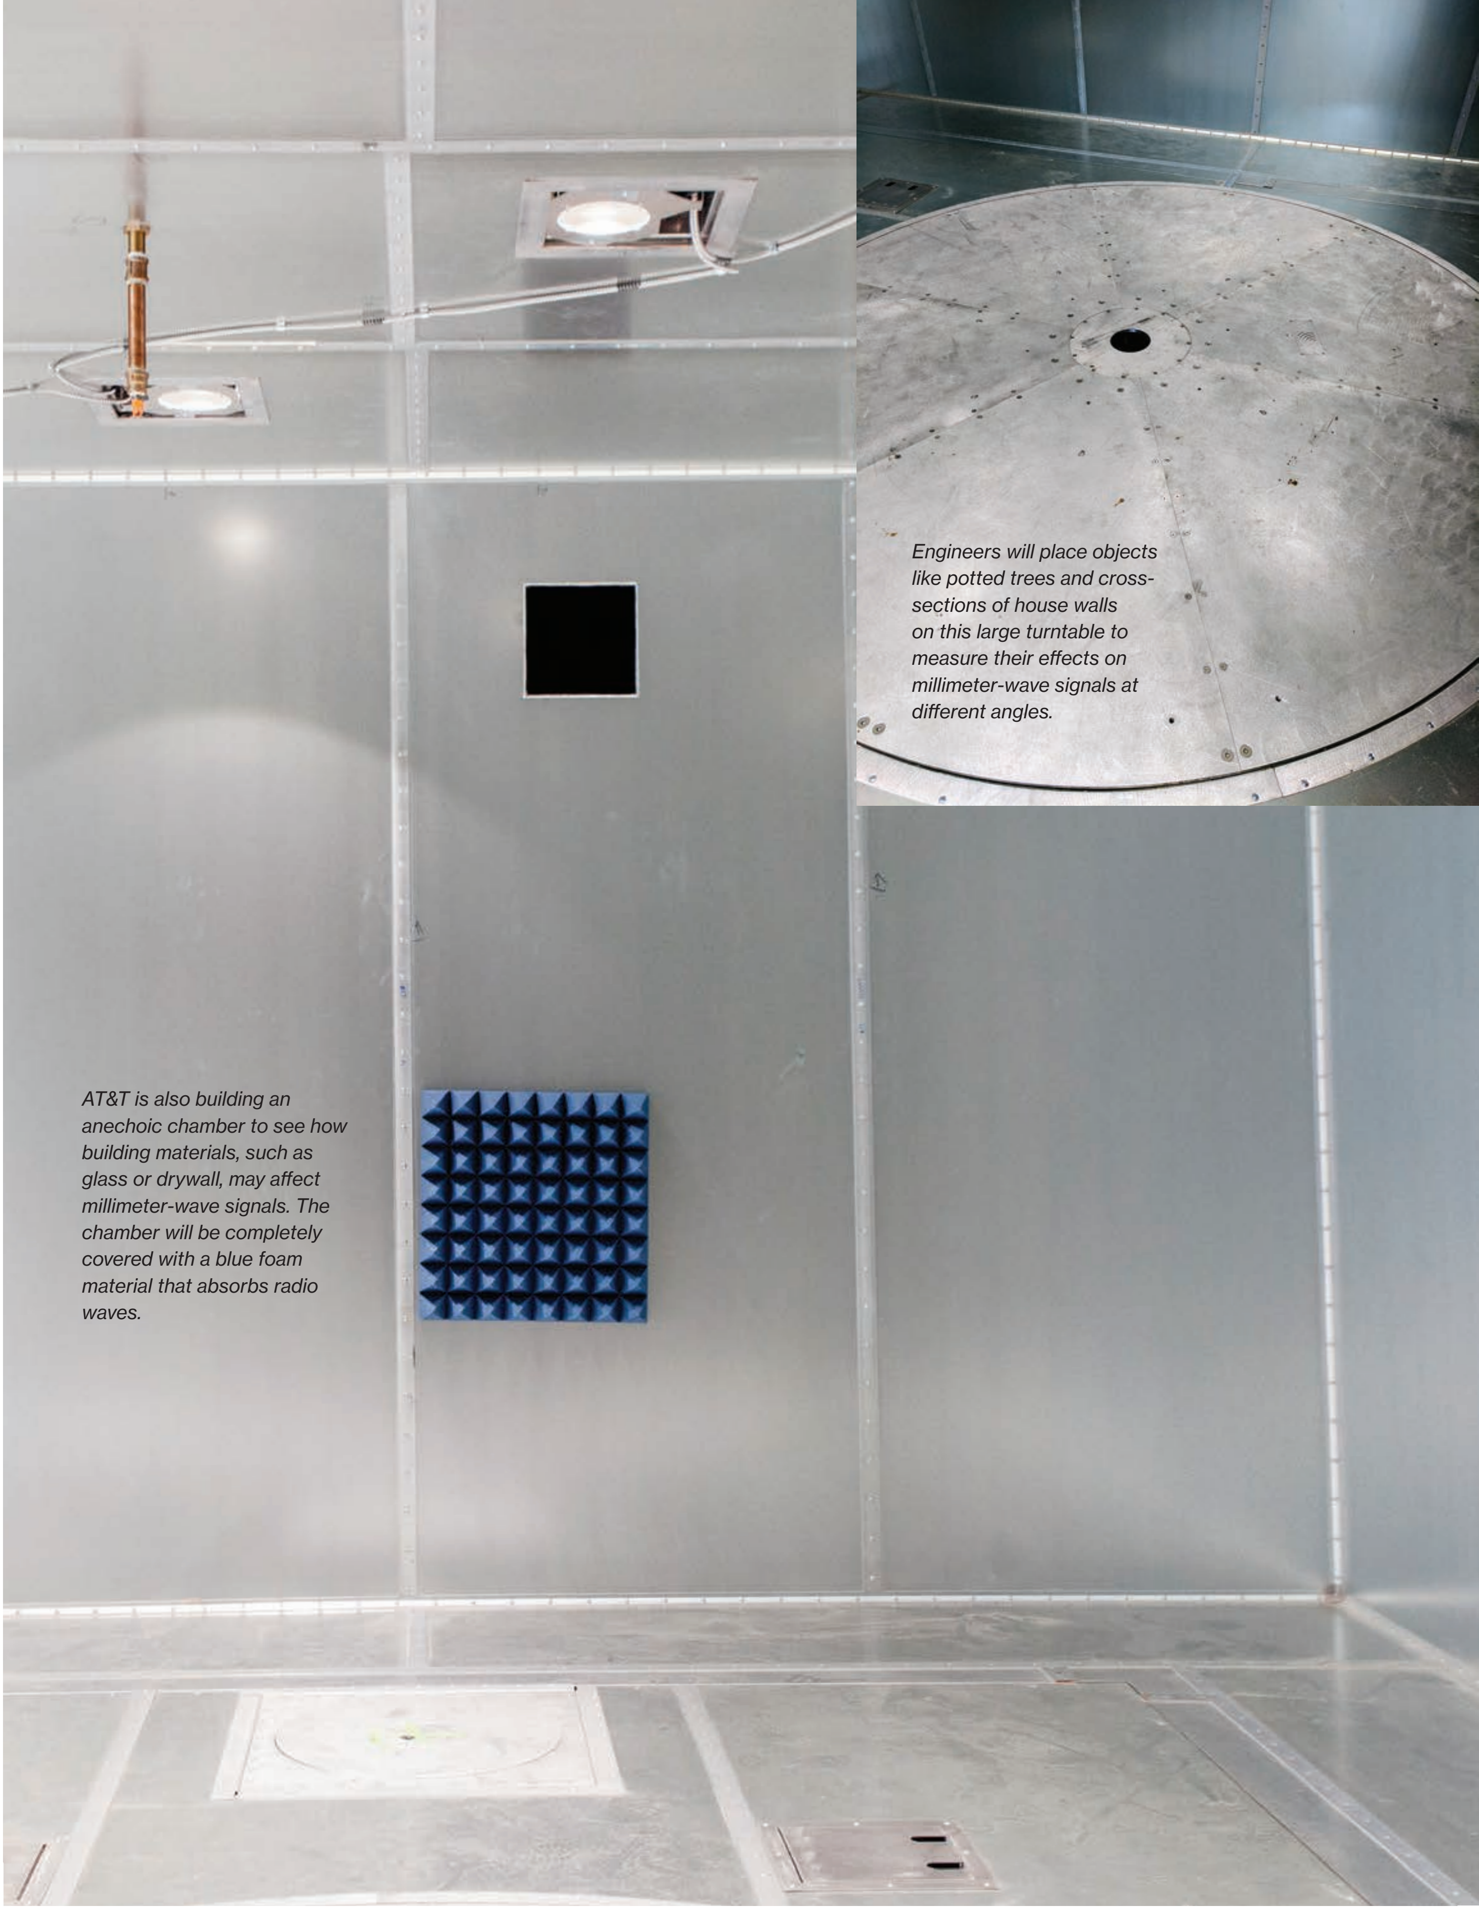

AT&T is also building an anechoic chamber to see how building materials, such as glass or drywall, may affect millimeter-wave signals. The chamber will be completely covered with a blue foam material that absorbs radio waves.

Engineers will place objects like potted trees and cross-sections of house walls on this large turntable to measure their effects on millimeter-wave signals at different angles.

bulky that it can be deployed outdoors for only a few hours at a time. Bennett and colleagues say that far more real-world data is needed to properly develop the technology, so they have created weatherproof radios the size of toaster ovens and installed them across AT&T's 260-acre campus in Middletown, New Jersey, which was once part of Bell Labs.

Since deploying the radios last September, the engineers have seen how tree leaves, heavy rain, and truck traffic all obstruct millimeter-wave signals to some extent. AT&T plans to share the information with the rest of the telecom industry to aid in the design of 5G technical specifications, base stations, modems, smartphone chips, and more. The new technology won't be commercially widespread until after 2020, but these small, homemade radios are a crucial step toward making it real.

Later this year, AT&T plans to install additional 5G measurement systems outdoors, and may use one of its testing vans to conduct tests outside the Middletown campus. The van's mast can extend up to 50 feet.

48 Years Ago

Taming the Weather

Nearly five decades ago it seemed as if we might gain mastery over the weather—but what of the risks?

"Most of man's inadvertent weather modification has resulted from pollution of the atmosphere by discharges from his technological establishment. Thermal pollution of the air over cities has caused them to become heat islands. Particulates and freezing nuclei have increased at all levels of the troposphere and have been related to increased cloudiness. The increase of atmospheric aerosols and particulates has been associated with an increase in atmospheric turbidity. Although direct measurements have not yet shown it, the greater cloudiness and turbidity threaten to reduce solar isolation. At the same time there has been a measurable rise in atmospheric carbon dioxide. Because this gas has a 'greenhouse' effect, several observers have forecast a general rise in the global air temperature of as much as 4 °C by the year 2000.

"Within the past decade man has had sufficient success with manipulating weather processes to encourage him to press on toward greater and more precise control. We need information that will help us to anticipate the nature and magnitude of biological changes that may result from a given weather modification.

"Some of these cause-and-effect relationships seem relatively simple: increased precipitation might favor tree growth but deep snow seriously impairs overwintering elk. But other biological effects may be remote from the target of the weather-modifying operation. Outbreaks of pests or of plant and animal diseases might materialize as a consequence of heavier amounts of rain and weaker animals. Weeds and undesirable animals might move into habitats made hospitable by weather modification. It is equally possible, however, that the altered weather might prove favorable for plants and animals desirable for man.

"Taming the weather can be a dangerous game. There may be gains for human welfare from achieving some control over the weather. At the same time, there are some imponderable ecological risks. Because man is as much a species of the biosphere as any other organism, he must critically evaluate these gains against the risks. He must undertake ecological studies that will allow him to anticipate the consequences and implications of conscious weather modification for the biosphere. In a real sense he holds his destiny in his own hands."

Adapted from "Weather Modification and the Biosphere," by Frederick Sargent, from the March 1969 issue of Technology Review.

MIT Technology Review (ISSN 1099-274X), May/June 2017 Issue, Reg. U.S. Patent Office, is published bimonthly by MIT Technology Review, 1 Main St. Suite 13, Cambridge, MA 02142-1517. Entire contents ©2017. The editors seek diverse views, and authors' opinions do not represent the official policies of their institutions or those of MIT. Periodicals postage paid at Boston, MA, and additional mailing offices. Postmaster: send address changes to MIT Technology Review, Subscriber Services, PO Box 5001, Big Sandy, TX 75755, or via the Internet at www.technologyreview.com/customerservice. Basic subscription rates: $41.94 per year within the United States; in all other countries, US$54. Publication Mail Agreement Number 40621028. Send undeliverable Canadian copies to PO Box 1051 Fort Erie, ON L2A 6C7. Printed in U.S.A. Audited by the Alliance for Audited Media


My Impossible:
After being diagnosed with Parkinson's, starting my own business didn't seem as scary.
Scott Kaufmann, Peer Aspect
Hiscox Customer
I'mpossible
Hiscox business insurance.
Helping the courageous overcome the impossible.
hiscox.com/impossible | 877-505-2048
HISCOX
encourage courage®


slack enterprise grid
Where work* happens.
* at large, fast-moving companies like yours
slack.com/enterprise